



Fermilab

FERMILAB-THESIS-1999-42

THE UNIVERSITY OF CHICAGO
THE ENRICO FERMI INSTITUTE

FIRST OBSERVATION OF THE RARE DECAY $K_L \rightarrow \pi^0 e^+ e^- \gamma$

A DISSERTATION SUBMITTED TO
THE FACULTY OF THE DIVISION OF THE PHYSICAL SCIENCES
IN CANDIDACY FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

DEPARTMENT OF PHYSICS

BY
GREGORY EDWIN GRAHAM

CHICAGO, ILLINOIS
AUGUST 1999

Copyright ©1999 by Gregory Edwin Graham

All rights reserved

ACKNOWLEDGMENTS

I would like to give thanks to everyone who has helped to make KTeV a reality. That the quality of the KTeV datasets is so high is due to the tireless efforts of many people. First I wish to acknowledge the KTeV spokespeople Y. Bob Hsiung, Yau Wah, Bruce Winstein, and Taku Yamanaka for their efforts in getting the KTeV experiments E832/E799 approved and realized. There were also many people in the trenches without whom KTeV could not have become the successful experiment that it is today. Ed Blucher, Elliott Cheu, Hogan Nguyen, Vivian O'Dell, Peter Shanahan, Sydney Taegar, Julie Whitmore, BoB Tschirhart, Tsuyoshi Nakaya, Rick Kessler, Rick Tesarek, Theo Alexopoulos, and John Belz are just a few of the names of collaborators who spent a seemingly endless supply of energy making sure KTeV was as successful an experiment as it was.

The analysis presented here also benefited from useful discussions with many people, of whom I would like to thank my advisor, Yau Wah, first and foremost. Also contributing were Rick Tesarek, Amit Lath, Taku Yamanaka, Eric Zimmerman, and all of the other “2 electron” meeting regulars.

The TRD system would not have been realized without the efforts of many people. Most importantly, the expertise and direction of John Krider led to a highly successful chamber design and construction plan, with important contributions by Nick Solomey. Much successful R&D was carried out and implemented through the efforts of John, Erik Ramberg, and Mike Daum. The electronics succeeded through the efforts of Wah, Harold Sanders, and John DuSatko at Chicago. Karen Kephart and the crew at Fermilab Lab 6 strung thousands of wires to desired tension and provided much technical advice. Elizabeth Pod contributed to the design and construction of the TRD radiators. I would also

like to thank Sasha Ledovskoy, Mike Arenton, and Cash McManus for their contributions to TRD debugging and commissioning. I would also like to thank Stephan Pordes for his dedicated support of the TRD project at Fermilab.

At this point, I would also like to thank two men who worked on the TRD but unfortunately did not live to see the results of their work : Dick Armstrong and Gene Beck. Dick Armstrong was a dedicated and hard working mechanical design engineer at the University of Chicago. He provided mechanical designs for most of the HEP projects for which the various Chicago groups were responsible. His contribution to the TRD system consisted of a system to align the TRD radiators to their chambers with much valuable advice throughout the project. Gene Beck was a hard working technician at Fermilab who really enjoyed his work and had the ability to brighten everyone's day with his jibes and jokes; a very difficult task at 7:00 AM. Gene's help during the early phases of TRD chamber construction was invaluable.

Finally, I would also like to give thanks to the graduate students of KTeV. They had no small role in each of its parts and phases. Below is a very incomplete listing. First I thank Eric Zimmerman for his work on the trigger and for holding his own with more senior physicists at bandwidth meetings. Val Prasad and Jim Graham expended much blood, sweat, and tears on the CsI array and the DC system respectively, Colin Bown on the HCC, Breese Quinn on the TRD system, and Peter Shawhan in general. I also thank Steve Bright and Val Prasad again for making KTeV more of a fun place to work. Peter Mikelsons, Eva Halkiadakis, Ashkan Alavi-Harati and many more also made many important contributions.

Lastly, I would like to give thanks to my wife Jennifer Pashup-Graham for her patience and support during the past eight years.

G. E. GRAHAM
Chicago, Illinois
July 1999

For my parents, Les and Lois.

ABSTRACT

A first observation of the rare decay $K_L \rightarrow \pi^0 e^+ e^- \gamma$ has been made by the KTeV E799 experiment at Fermilab. Based on a sample of 45 events with an estimated background of 3.5 ± 1.3 events, we obtain $BR(K_L \rightarrow \pi^0 e^+ e^- \gamma) = (2.42 \pm 0.38 \pm 0.11) \times 10^{-8}$ where the first error is statistical and the second systematic. This result is in agreement with recent calculations using Chiral Perturbation Theory at $O(p^6)$.

TABLE OF CONTENTS

ACKNOWLEDGMENTS	iii
ABSTRACT	vi
LIST OF TABLES	xi
LIST OF ILLUSTRATIONS	xiv

Chapter

1	Introduction to $K_L \rightarrow \pi^0 e^+ e^- \gamma$	1
1.1	Introduction to $K_L \rightarrow \pi^0 e^+ e^- \gamma$	2
1.2	Why is $K_L \rightarrow \pi^0 e^+ e^-$ a CP -Violating Decay ?	3
1.2.1	CP Symmetry and the Neutral Kaons	4
1.2.2	CP Conserving Contribution to $K_L \rightarrow \pi^0 e^+ e^-$	6
1.3	$K_L \rightarrow \pi^0 e^+ e^- \gamma$ and Its Relationship to $K_L \rightarrow \pi^0 e^+ e^-$	9
1.4	Experimental Technique	10
2	Introduction to the Beam and Detector	12
2.1	The KTeV Beam	12
2.1.1	The KTeV Secondary Beam	13
2.1.2	The Timing of the KTeV Beam	16
2.1.3	Final Characteristics of the KTeV Beam	17
2.2	The KTeV Detector	19
2.2.1	The KTeV Vacuum Decay Tank and Ring Counter Vetoes	20
2.2.2	The Drift Chamber Spectrometer System	22
2.2.3	The Transition Radiation Detectors (TRDS)	29
2.2.4	The Trigger Hodoscopes	30
2.2.5	The Cesium Iodide Calorimeter	30
2.2.6	The Lead Wall and Hadron Anti	36
2.2.7	And the Rest	38

3	Locating the Interesting Events	39
3.1	The Trigger	39
3.1.1	The L1 Trigger	41
3.1.2	The L2 Trigger	42
3.1.3	The L3 Filter	45
3.1.4	Other Triggers	45
3.2	The Data Acquisition System	46
3.2.1	The CsI Readout and Digital Pipeline	46
3.2.2	Other Readout	47
3.2.3	Streams, Planes, and Event Building	48
3.2.4	The Slow DAQ : Monitoring and Control	49
3.3	Further Refinements of the Data Set	50
3.4	Conclusion	53
4	The Transition Radiation Detector in Detail	56
4.1	Overview of the KTeV TRD System	57
4.2	Optimization of the TRD System with Simulation	57
4.3	The Radiator	60
4.4	The X-Ray Detector	64
4.5	Testing and Commissioning	69
4.6	The TRD Electronics	71
4.7	CAEN High Voltage Control and Monitoring	72
4.7.1	CAEN Server Program	74
4.7.2	CAEN Clients	76
4.8	Gas System	79
4.9	TRD Performance	82
5	Elementary Event Reconstruction	87
5.1	The Database	87
5.2	Event Reconstruction	88
5.2.1	A Sample KTEVANA Job	88
5.2.2	Clustering	90
5.2.3	Tracking and Vertexing	97
6	Monte Carlo Simulation	102
6.1	General Aspects of the Monte Carlo	103
6.1.1	The Life and Times of a Monte Carlo Event	103
6.1.2	The Momentum Distribution and K0K0BAR Mix	106
6.1.3	Particle Tracing and Detector Response	109
6.1.4	Z Position of Decay Vertex	109
6.1.5	Particle Tracing	109
6.1.6	Photon Vetoes	111
6.1.7	Drift Chambers	111
6.1.8	Hodoscopes	112
6.1.9	CsI Calorimeter	112

6.1.10	Accidentals Simulation	116
6.2	Individual Decay Generators	116
6.2.1	Monte Carlo Simulation of $K_L \rightarrow 2\pi^0$	116
6.2.2	π^0 Dalitz Decay	117
6.2.3	Monte Carlo Simulation of $K_L \rightarrow \pi^0 e^+ e^- \gamma$	118
6.2.4	Monte Carlo Simulation of $K_L \rightarrow e^+ e^- \gamma \gamma$	121
6.2.5	Monte Carlo Simulation of $K_L \rightarrow 3\pi_D^0$	122
7	The Normalization Mode Analysis	123
7.1	Cuts To Isolate $K_L \rightarrow 2\pi^0, \pi^0 \rightarrow e^+ e^- \gamma$	126
7.1.1	K_S Backgrounds	128
7.1.2	$K_L \rightarrow 3\pi_D^0$ Background	132
7.2	Cuts To Improve Data Quality	134
7.3	Flux Results and Systematic Error	134
7.3.1	Data - Monte Carlo Comparisons	135
7.3.2	Plots Over Which I Lose Sleep	141
7.4	Conclusion	151
8	Analysis of $K_L \rightarrow \pi^0 e^+ e^- \gamma$	153
8.1	Analysis Cuts for $K_L \rightarrow \pi^0 e^+ e^- \gamma$	155
8.1.1	Cuts to Identify $K_L \rightarrow \pi^0 e^+ e^- \gamma$	156
8.1.2	Cuts to Remove Backgrounds	157
8.2	Results on $K_L \rightarrow \pi^0 e^+ e^- \gamma$	178
8.2.1	Summary of Calculated Backgrounds to $K_L \rightarrow \pi^0 e^+ e^- \gamma$	178
8.2.2	Acceptance of $K_L \rightarrow \pi^0 e^+ e^- \gamma$	179
8.2.3	Systematic Error	179
8.2.4	Tying It All Together	180
8.2.5	First Observation pf $K_L \rightarrow \pi^0 e^+ e^- \gamma$	181
8.2.6	$K_L \rightarrow \pi^0 e^+ e^- \gamma$ as a background to $K_L \rightarrow \pi^0 e^+ e^-$	182
9	Conclusion	185
A	Note on Error Estimates with Weighted Monte Carlo	187
A.1	Working with Weighted Monte Carlo	188
A.1.1	Errors with Weighted Monte Carlo	188
A.1.2	Estimate of the Number of Degrees of Freedom	190
A.2	Test of the Error Estimates	191
A.2.1	Errors with Weighted KTeV Monte Carlo	191
A.2.2	Results with Simple Random Number Monte Carlo	192
A.3	Conclusion	194
A.4	A Small Example	195

B	Confidence Intervals With Background Error :	
An	Extension to the Method of Feldman and Cousins	196
B.1	The Method of Feldman and Cousins in Brief	197
B.2	An Extension to the Method of Feldman and Cousins	199
B.3	Short Table of Confidence Intervals With Errors in Background	201
C	Multi Particle Kinematics in Two E-Z Pieces	203
REFERENCES	208

LIST OF TABLES

1	A Table showing different ways of thinking about the neutral kaon.	6
2	Z positions and aperture information for important elements in the KTeV secondary beam system. Defining apertures are listed in bold	15
3	Relative populations of various particle species surviving in the E799 beam.	19
4	Z positions and relevant apertures for elements in and around the KTeV main decay volume. The limiting apertures are in bold	21
5	KTeV Spectrometer Z positions and Apertures. Only the inner apertures are given. The Z positions are at the upstream face of the detector.	28
6	Breakdown of the radiation length of the one of the KTeV TRD chambers. The total is calculated for the in-acceptance part of the chamber. In the regions intersecting the neutral beams, there is no radiator material.	29
7	Z positions and Apertures for the trigger hodoscopes and the TRD system. None of these apertures is limiting.	30
8	Z positions and apertures for detector elements around the CsI. The limiting apertures are in bold	32
9	Detector elements downstream of the CsI with Z positions. There were no limiting apertures here. Only the HA is relevant to this analysis through the trigger, and MU2 likewise for part of the run.	37
10	Bad Spills for the Winter and Summer E799 runs. The numbers in the table represent fully the number of fully reconstructed $K_L \rightarrow 2\pi_D^0$ decays with the bad spill reasons shown at the left. The data in the table are compiled post-run 8245 only. Before run 8245, every event was flagged by problem 17 due to HCC cable swap. There were 2230 reconstructible $K_L \rightarrow 2\pi_D^0$ events with this problem representing about 7.5% of the winter data. Furthermore, we do not cut on problem 3 in the summer because for the majority of cases, the presence of problem 3 reflects an error in the database rather than a true problem with the DPMTs.	54
11	Permeability of several gases in the KTeV TRD system across 2 mil mylar windows, aluminized or not aluminized. The units are in $\frac{cc}{day \times 100in^2 \times atm}$. (Based upon a table given in [KTeV 154].)	65

12	The X-ray absorption lengths, L_x , of various materials are presented with material densities. The X-ray absorption lengths are calculated from fits to data in [Storm 70].	66
13	The parameters found by Malensek to fit the production of charged kaons by protons on a Be target. See equation 6.5.	108
14	Radiation lengths of selected components following the KTeV decay volume. Radiation lengths are modeled as planes that contain the compressed amount of material from the surrounding regions. The helium bags are modeled with four such planes to better model the continuous nature of the physical regions. <i>Courtesy of E.D. Zimmerman.</i>	110
15	Mass resolutions for the various π^0 masses calculated in the normalization mode analysis. The mass peaks in each case were fit to a single Gaussian in a close range about each peak. In all cases, the fit χ^2 was between 1.5 and 2.0, and the quoted errors are the fit errors. The quoted resolution is equal to the sigma of the fit Gaussian.	126
16	Mass resolutions for $e^+e^-\gamma\gamma\gamma$ mass calculated in the normalization mode analysis. The mass peaks in each case were fit to a single Gaussian in a close range about each peak. In all cases, the fit χ^2 was between 1.5 and 2.0, and the quoted errors are the fit errors. The quoted resolution is equal to the sigma of the fit Gaussian.	127
17	The acceptances and numbers of observed $K_L \rightarrow 2\pi_D^0$ events in E799. The winter run is split at run 8245 to take account of a slight acceptance loss before run 8245 due to a swapped HCC cable.	135
18	The E799 K_L flux. The winter run is split at run 8245 to take account of a slight acceptance loss before run 8245 due to a swapped HCC cable. The quoted errors are statistical only.	136
19	A table of systematic errors in the normalization mode analysis in order of importance. The total estimated systematic error in the flux measurement in 1.7%.	142
20	The X and Y positions of the DC regions intersecting the beam in the winter and the summer for drift chambers 1 and 2.	145
21	Mass resolutions for the various π^0 masses calculated in the normalization mode analysis and the $K_L \rightarrow \pi^0 e^+ e^- \gamma$ analysis. The mass peaks in each case were fit to a single Gaussian in a close range about each peak. In all cases, the fit χ^2 was between 1.5 and 2.0, and the quoted errors are the fit errors. The quoted resolution is equal to the sigma of the fit Gaussian.	156
22	Mass resolutions for $e^+e^-\gamma\gamma\gamma$ mass calculated in the normalization mode analysis and the $K_L \rightarrow \pi^0 e^+ e^- \gamma$ analysis. The mass peaks in each case were fit to a single Gaussian in a close range about each peak. In all cases, the fit χ^2 was between 1.5 and 2.0, and the quoted errors are the fit errors. The quoted resolution is equal to the sigma of the fit Gaussian.	157
23	The relevant background contributions to the final $K_L \rightarrow \pi^0 e^+ e^- \gamma$ sample. Backgrounds from double conversion events and four track events are very small and are not included.	178
24	Systematic errors in the measurement of $K_L \rightarrow \pi^0 e^+ e^- \gamma$	180

25	Cut flow table for the $K_L \rightarrow \pi^0 e^+ e^- \gamma$ analysis. The normalization mode cuts are discussed in chapter 7. The “ $2\pi_D^0$ Killer” refers to the mass cuts used against $K_L \rightarrow 2\pi_D^0$ in all photon pairing hypotheses. The “ $3\pi_D^0$ Killer” refers to the kinematical cut against $K_L \rightarrow 3\pi_D^0$. The cut values are described in the text. The first column lists the number of events surviving in the data. The second column lists the acceptance (combined winter and summer numbers.) The third and final columns detail the normalized expected backgrounds from the two major background sources, $K_L \rightarrow 2\pi_D^0$ and $K_L \rightarrow 3\pi_D^0$. These columns are normalized so that the events that pass all cuts correspond to the final estimates of background from those sources.	182
26	Errors for calculated acceptances using KTeV $K_L \rightarrow e^+ e^- \gamma \gamma$ Monte Carlo. 50 trial runs were performed to measure the acceptance. The errors were calculated for each trial run and compared to the observed scatter of individual means A_j about the global mean A_{tot} . Listed in the table are the error expression used to calculate the errors in each trial run, the simple average error calculated in the trial runs, and the χ^2 statistic for the 50 trial runs compared to the global average computed following equation A.14. Also listed at the end of the table are the average sizes of the corresponding terms in expression A.5. In the case of KTeV $K_L \rightarrow e^+ e^- \gamma \gamma$ Monte Carlo, the last two terms of expression A.5 are much smaller than the first term.	192
27	Errors for calculated acceptances using simple randomly generated weights and acceptances. Each number in the table represents the χ^2 per degree of freedom fit of the averages of a set of 100 trials of 5000 events each to the global average. A low χ^2 indicates that the errors are overestimated in that set of trials. The χ^2 was calculated using both equation A.5 and equation A.8 for three different cases : acceptance is uncorrelated with weight, acceptance is correlated with weight, and acceptance is anti-correlated with weight. The normalization of the acceptance is then varied so that the average acceptance varies from 1% to 75%.	193
28	A small ”pencil and envelope” sample of events. The events are characterized solely by their weights and their acceptance bits.	195
29	90% confidence Intervals calculated using the method of Feldman and Cousins with exact errors and with 10% and 20% Gaussian errors in the background. The cases with Gaussian error were each calculated on a grid of 11 points $\pm 5\sigma$ about the stated mean background. Most limits do not change, but many change by up to a few percent. . .	202

LIST OF ILLUSTRATIONS

1	The diagram representing CP conserving $\pi\gamma\gamma$ intermediate state to $K_L \rightarrow \pi^0 e^+ e^-$	2
2	Generalized diagram for the rare decay $K_L \rightarrow \pi^0 e^+ e^- \gamma$. If we compare this figure with figure 1, we see that the $K_L \rightarrow \pi^0 e^+ e^- \gamma$ process has one less internal electro-magnetic coupling. This means that the rate for $K_L \rightarrow \pi^0 e^+ e^- \gamma$ should be enhanced relative to the CP conserving contribution to $K_L \rightarrow \pi^0 e^+ e^-$ by a factor of the electro-magnetic coupling, α	3
3	The dependence of BR on a_v . <i>Reproduced courtesy of F. Gabbiani.</i>	8
4	Diagram detailing the $K \rightarrow 3\pi$ contribution in chiral perturbation theory to $K_L \rightarrow \pi^0 e^+ e^- \gamma$. (For $K_L \rightarrow \pi^0 e^+ e^- \gamma$, one of the final photons is off shell.)	9
5	Two diagrams detailing vector meson exchange contributions to $K_L \rightarrow \pi^0 e^+ e^- \gamma$. These contributions only appear at $O(p^6)$ in the calculation. (For $K_L \rightarrow \pi^0 e^+ e^- \gamma$, one of the final photons is off shell.)	10
6	A sketch of the fixed target area at Fermi National Accelerator Laboratory. The KTeV experiment is located on the neutrino-muon (NM) beam-line. Note that the map given in the Fermilab telephone directory is wrong. Not shown : Rosati's.	13
7	KTeV beam elements. Kaons are produced by 800 GeV/c protons on a BeO target at Z=0 m. Charged particles are removed from the beam at various points by sweeping magnets while the neutral beam is shaped by a set of collimators. A lead absorber in the beam-line serves to convert photons in the beam to $e^+ e^-$ pairs, which are later swept out. The common absorber was not used in E799.	14
8	Simulated K_L momentum in E799. The top plot shows the raw generated kaon momentum spectrum derived from the Malensek spectrum. The bottom plot shows the K_L momentum spectrum for K_L decays which satisfy the triggering requirements for the analysis considered in this thesis. Decays of K_L particles below about 20 GeV/c fail to satisfy triggering requirements because they tend to decay upstream and have decay products that escape the detector.	18
9	The KTeV detector. The scale in the Z direction is compressed with respect to the X direction by a factor of about 11.	20

10	The drift chamber geometry looking along the positive Y axis. The open boxes represent the field shaping wires and the crosses represent the sense wires. (The solid lines represent field shaping wires and the dashed lines represent the sense plane in the rotated view.) The upstream pair of sense planes consists of wires oriented along the Y direction to provide a measurement of X position. The downstream pair of sense planes consists of the planes along the X direction for a measurement of Y position. Adjacent sense wires are 12.7 mm apart. The two sense planes in each view are staggered to allow for unambiguous position measurement.	23
11	The arrival times of signals in DC2 during run 8397. Operated in common-stop mode, the hits close to a wire start the TDC counting earlier than other hits and contribute to the “bluff” near 350 ns. Later hits trail off to the left.	25
12	The sum of distances, or <i>sod</i> , distribution for $K_L \rightarrow 2\pi_D^0$ events in DC2 in the E799 winter run. The mean of the sod distribution corresponds to $\frac{1}{2}$ a cell width in the DC, as seen in figure 10.	26
13	The V and V' trigger hodoscope planes. The paddles shown are of different widths so that there are no overlapping cracks when one plane is overlaid on the other.	31
14	The CsI calorimeter was constructed from 3100 crystals of pure CsI in a square array with two beamholes in the center to allow the neutral beam to pass through. The size of the array is 1.9 m \times 1.9 m. Two 15 cm square beamholes exist at (X,Y) = (± 15.0 cm,0.0cm) to allow for passage of the neutral beam. A 1.2 m decorated hardboiled egg is shown for scale.	33
15	A diagram showing possible pattern inputs into the HCC and their weights. In the HCC, all 2 \times 2 patterns of crystals are processed according to which crystals are above threshold. The sum of the weights taken over the entire array is 4 times the number of clusters.	35
16	In this closeup view of the CsI beamholes, one can get a glimpse of the wily Collar Anti (CA). This beast lives in front of the beam holes of the CsI, overlapping the innermost layer of crystals by 1.5 cm. The CA guards the hole edges in the CsI and vetoes events which are too close.	36
17	A drawing of the HA counter bank. The HA sat immediately downstream of the Pb wall and was used as a veto on the detection of hadronic showers at Level 1.	37
18	The selection of in-time hit pairs by the Banana modules. The figure on the left shows the selected in-time hit pairs which are shaped like a banana. The upper triangle above the banana consists of low-sod pairs which were often reconstructible offline; therefore these were also accepted at trigger level. The figure on the right shows the out of time bananas. The bananas could either be early (unripe) or late (spoiled.) Also, events with unphysical sods may not contribute to any of the bananas.	44

19	A depiction of the DAQ system. Event fragments traveled from the front end crates in 6 streams, each of which had 40 Mb/sec bandwidth. Fragments from a contemporaneous event were held in a single DAQ plane and analyzed by the L3 Filter. On a L3 accept, the event was to a contiguous block of memory and out to tape. The monitoring plane was the destination plane for monitoring events and calibration triggers. Calibration trigger data was written to tape from plane 3.	48
20	Bad spill reasons in E799 from reconstructed $K_L \rightarrow 2\pi_D^0$ events. Bad spill reasons 24,26, and 28 correspond to TRD problems and are not checked in this analysis. . . .	55
21	The TRD system is depicted along the beam axis. All 8 TRDs lie between the last drift chamber, DC4, and the trigger hodoscope planes. A spare TRD module, not shown, was also in place during the summer run immediately downstream of TRD8.	58
22	The TRDs are shown in cross section to the beam. During E832, the TRDs could be rolled out of acceptance. Flexible cable trays were used so that no special cabling operations were necessary during rollout or roll-in.	58
23	The conceptual design of a KTeV TRD module. Xrays generated in the radiator must pass through a double window buffer volume containing an 80/20 mixture of C2F6/CO2. The width of this region is up to 1 cm. The Xrays are then detected in an active region containing an 80/20 mixture of Xe/CO2. The active region contains the MWPC, consisting of two anode planes and three alternating cathode planes. The anode - cathode gap is 6mm, and the cathode wire pitch is 5 mm. The anode wire pitch is twice the cathode wire pitch, and the two anode planes are staggered from each other. . . .	59
24	For a constant amount of radiator material, the number of TRD modules in the system was varied and performance modeled by monte carlo. The estimated pion rejection plateaus around N=8. Below this value, the pion rejection falls rapidly. The radiator thickness at N=8 corresponds closely to a single X-ray absorption length at 10 keV. For smaller values of N, the radiators are thicker than this value, and X-rays are lost to radiator self-absorption.	61
25	The radiator frame is shown in cross section to the beam and along the beam direction. The frame itself is made from a light core of plastic honeycomb clad with eighth inch thick aluminum plate. The frame was stretched with a 2 mil aluminized mylar window to impede movement of water vapor across the boundary. The radiator stack was made from polypropylene felt blankets and were hung from pegs at the top of the frame. In addition, the radiator was glued down at the edges. Finally, holes were cut in the radiator in the beam regions to allow passage of the two KTeV neutral beams. The aluminum coating was etched away from the mylar around the beam holes so that the final position of the beam holes could be monitored after the radiators were mounted to the TRD chambers.	63

- 26 A diagram showing the components of a TRD frame. There were three layers in each lamination. The middle layer was a single piece steel core depicted at left. The outer layers were each four pieces; G10 on top and bottom, and aluminum on the sides. The medium separating the lamination layers was provided by a thin sponge material and glue. When pressed together between two externally aligned flat surfaces and allowed to dry, a flat wire plane resulted. The front of the wire plane had G10 pieces clad with copper traces to provide mount points for wires to be stretched later. 64
- 27 The dual electroplating boat used to deaden wires in the KTeV TRD beam regions. The dead regions were large enough to contain the KTeV beams (E799 summer large beams) within a 3mm tolerance for all TRDs. 68
- 28 A diagram showing the major steps in the construction of the TRD chambers. The primary mechanical support comes from an aluminum strongback, which is 6" aluminum square hollow channel. An aluminum frame on which is stretched the two aluminized mylar windows comprises the C2F6 buffer volume. The steel-G10 laminated wire frames, both holding the wires and providing the anode-cathode separation distance, are stacked next. Finally, the rear window is added followed by an additional aluminum tube support. These elements are sealed with RTV sealer and screwed down to uniform tolerances. 69
- 29 The CAEN HV system as used in KTeV. Remote communication with the CAEN SY527 mainframes in the pit took place through a C117B CAMAC interface. An emergency kill bypass through direct front panel interface also existed but was never used. . . . 73
- 30 The CAEN server process. The CAEN server listened on open message queues. The "G" (General) queue handles run-of-the-mill requests for HV monitoring and control. The "P" (Priority) queue handles emergency requests for voltage shutdown. Since the "P" queue is checked before the "G" queue during each processing cycle, messages on "P" cannot be locked out by multiple messages on "G". Clients supply the server with a return queue for data transfer if needed. 75
- 31 The active HV control system receives input from a barometer. During E799, the barometer was checked every 15 minutes. A gain shift was calculated based upon the shift in pressure from the last check. Using a calibrated voltage - gain curve for each TRD plane, a compensating voltage shift was calculated. The active HV process then adjusted the TRD high voltages automatically through the CAEN server. 79
- 32 The performance of the active HV system during the part of the E799 winter run is plotted above for TRD1. Before run 8200, the TRDs were not running reliably for various reasons. After run 8200, the TRDs were online and the active HV system was running. For a brief period around run 8480, the active HV failed. Interestingly, this active HV failure corresponds to a sharp bump in the plot. 80
- 33 The N2 remover designed by J. Krider at Fermilab. The system depends upon the difference of permeation rates through silicone tubing for N2 and Xe. Xe diffuses through more quickly and is therefore trapped in the canister. The Xe depleted gas is exhausted. 81

34	TRD ADC spectra are shown from the front plane (top) and the back plane (top) of TRD4. The events shown are identified Ke_3 events. The electrons have a harder spectrum than the pions due to the addition of transition radiation.	83
35	For two ADC measurements x_1 and x_2 , the probabilities that a pion generates a signal greater than these values are P_1 and P_2 respectively. The total probability is the product of the two individual probabilities, which generates a hyperbola in the probability space as shown above. A confidence level can be given as the integral of the probability space below the hyperbola.	84
36	The pion probability is shown for a sample of Ke_3 electrons overlaid with Ke_3 pions. A cut in pion probability is used to reject pions.	85
37	The single track pion rejection performance at 90% electron efficiency is shown in all of the E799 KTeV TRD calibration ranges. The winter run comprises calibration ranges 08a through 33b. The summer run comprises calibration ranges 37b through 55a. <i>Plot courtesy of L. Bellantoni.</i>	86
38	The energy calibration in a selected run range of E799. The calibration uses the property that electro-magnetic showers are completely contained in the CsI, so that the mean ratio of energy to momentum, or E/p , for electrons is 1.0. The figure at left shows the E/p distribution for electrons from Ke_3 decays before the offline energy calibration. The figure at right shows the distribution after offline calibration for the same set of data. The mean E/p after calibration is 1.0006 by Gaussian fit, and the resolution is 0.99%.	91
39	The threshold correction is depicted for a small crystal that is located 3 blocks above and three blocks to the side of the seed block. The mean energy as a fraction of the applied threshold is plotted versus cluster energy for under threshold blocks. The data shown are electrons from Ke_3 events in run 9897, a special run with no CsI readout threshold.	93
40	Mean E/p versus run number for E799 runs in the “kqkq04” range, an arbitrary run range. The top plot shows E/p in clusters consisting entirely of large crystals and the bottom shows E/p in clusters consisting entirely of small crystals. (The relevant calibration ranges are (a) 8358 to 8377, (b) 8378 to 8407, (c) 8408 to 8451, and (d) 8452 to 8482.) The run to run variations are much greater in the small crystal region in the center of the CsI array where radiation damage is greater. The calibration ranges were chosen to limit the overall variation to 1% in any range, and corrections were later applied based on plots like these within each range.	94
41	Mean E/p is plotted versus electron momentum for electrons from Ke_3 events in clusters consisting of large crystals only. This is the basis for the final non-linearity fudge for the large crystals.	95
42	Mean E/p is plotted versus electron momentum for electrons from Ke_3 events in clusters consisting of small crystals only. This is the basis for the final non-linearity fudge for the small crystals.	96

43	Good SOD, Low SOD, High SOD, No SOD. The different kinds of SODs are shown. Typically, a single track traverses a pair of drift cells so that the sum of drift distances add up to $\frac{1}{2}$ of a cell spacing. However, if two tracks or a single track with a delta ray are present, a low SOD pair will result. The high SOD pair occurs when a track passes very close to a wire as explained in the text. Finally, if a wire is inefficient, an isolated hit can result.	97
44	A comparison of SOD distributions between E799/E832 and the previous experiment E773. Both experiments used the same chambers and wire geometry.	98
45	The Malensek production cross sections used to pick initial momentum and direction of generated kaons in the KTeV Monte Carlo simulation for the “K0K0BAR” option. The distributions also reflect weighting from the probability of decay in the specified momentum ranges and Z ranges. Thus, the survival of high momentum K_S into our decay volume is represented by small enhancements in both the K^0 and \bar{K}^0 plots. . .	107
46	A relationship between the neutral kaon production cross sections and the measured charged kaon cross sections can be estimated by simple quark counting arguments in the beam proton. The beam proton has two u quarks and one d quark in the valence. The sea quark content relative to the valence quarks can be parameterized by x for the $u\bar{u}$ or $d\bar{d}$, and y for the $s\bar{s}$, as shown in part (a) of the figure. Using the quark assignments for the charged and neutral kaons, one can deduce the relationships shown in part (b) of the figure and those in equation 6.4.	107
47	The position bins used to store GEANT generated electro-magnetic showers for the CsI simulation. Shower lookups for regions of the crystal outside of the indicated region can be related to the defined bins by symmetry.	114
48	The decay distribution in Q as found by the KTeV Monte Carlo. This decay generator is based upon that used in [Gabbiani 97].	119
49	The decay distribution in Y as found by the KTeV Monte Carlo. This decay generator is based upon that used in [Gabbiani 97].	120
50	The decay distribution in Z as found by the KTeV Monte Carlo. This decay generator is based upon that used in [Gabbiani 97].	121
51	The $K_L \rightarrow 2\pi_D^0$ decay final state. Two photons come from a π^0 , while the remaining photon and e^+e^- pair also come from a π^0 which decays via the Dalitz decay mode. .	124
52	The $e^+e^-\gamma\gamma$ mass resolutions using charged and neutral vertex positions. Since the neutral vertex calculation implicitly constrains the $\gamma\gamma$ mass to be the π^0 mass, improved resolution results.	125
53	The π^0 and π_D^0 mass peaks in the winter E799 dataset for the $K_L \rightarrow 2\pi_D^0$ sample after all cuts with monte carlo overlay. The resolutions match very well, and the means between data and monte carlo agree to within 0.15%.	127

54	Collimator Scatters. The top plot shows the log distribution of transverse momentum for E799 $K_L \rightarrow 2\pi_D^0$ data after all cuts. The filled circle overlay is the $K_L \rightarrow 2\pi_D^0$ monte carlo prediction with no collimator scatters, and the filled triangle overlay is the collimator scatter ONLY monte carlo prediction. Kaons scattering in the collimator tend to regenerate to K_S so that the fraction of collimator scatters into 2π is relatively high. The cut shown is the final analysis cut in P_T to remove the collimator scatters. The bottom plot shows the extrapolated XY position of the kaon in the collimator plane for events which fail the P_T cut. The resulting plot is aptly called the “hotlips” plot. . . .	129
55	The top plot shows the distribution of the final E799 $K_L \rightarrow 2\pi_D^0$ data in vertex Z position and kaon momentum. The region at high momentum closer to the target is populated by decays of K_S that survive from the target. The bottom plot shows the Z/P ratio for $K_L \rightarrow 2\pi_D^0$ monte carlo events. The K_S survivors from the target dominate the region $Z/P < 0.75$. The filled circle overlay are the same monte carlo events weighted to be K_S , and the filled triangle overlay are the same events weighted to be K_L according to equation 7.4.	131
56	$e^+e^-\gamma\gamma\gamma$ invariant mass is plotted versus log of $ P_T $. $K_L \rightarrow 3\pi_D^0$ background as well as collimator scatters are well excluded by the indicated cuts in mass and P_T . Also, a popular corporate logo.	132
57	The $e^+e^-\gamma\gamma\gamma$ final state mass for the $K_L \rightarrow 2\pi_D^0$ normalization mode after all cuts for the winter E799 dataset with monte carlo overlay. The background at low invariant mass below the kaon peak is from $K_L \rightarrow 3\pi_D^0$ fragments. In the mass window between 487 and 507 MeV/ c^2 , there are estimated to be about 10 $K_L \rightarrow 3\pi_D^0$ events based upon a fit of the background shape in the low mass sideband.	133
58	X position illumination of tracks at DC1 in the summer and the winter with monte carlo overlay.	136
59	Y position illumination of tracks at DC1 in the summer and the winter with monte carlo overlay.	137
60	As a measure of the beam illumination, the ratio of vertex X and Y positions to the vertex Z position are plotted for both the summer and winter datasets with monte carlo overlay. The slope in the V_Y/V_Z distribution is due to the targeting angle.	138
61	Cluster X and Y position illumination at the CsI for both the summer and winter datasets with monte carlo overlay.	139
62	Kaon momentum distribution in the winter and summer.	140
63	The distribution of number of beam hits per event in DC1 and DC2 with monte carlo overlay. The top pair of plots show the overlay for E799 winter data, and no evidence for beam region inefficiency is found. The bottom pair of plots show the overlay for the E799 summer data, and a small constant inefficiency per beam hit is seen.	141

- 64 The Z slope problem in $K_L \rightarrow 3\pi_D^0$. In order to plot the dependence of the Z slope problem on opening angle, a 2 dimensional scatterplot (upper left) is filled with data from the plentiful $K_L \rightarrow 3\pi_D^0$ decay mode. The Z vertex position is plotted versus opening angle between the two electron tracks. Note that for this mode, like all π_D^0 modes, the plot is dominated by low opening angles. The scatterplot is sliced in opening angle, and the vertex Z position distribution is compared between data and MC for each slice (upper right). The ratio plot in each slice is fit to a line. Finally, the slope is plotted versus the opening angle (bottom plot.) ($K_L \rightarrow 3\pi_D^0$ data provided courtesy of Pat Toale.) 143
- 65 The SOD (sum of distances) distribution for plane pair DC2X in the winter E799 data from the $K_L \rightarrow 2\pi_D^0$ sample. The right side shows the data with monte carlo overlay, and the left side shows the ratio plot. The ratio plot shows the mismatch of the widths as a parabolic curve. The means of the two distributions is also offset from one another. 144
- 66 The cartoon picture above shows the definition of drift chamber regions for a drift chamber. If the plane pair is in the X view, then there are five regions defined by the two strips containing beam holes and the three surrounding regions. In the Y view, there is only one strip containing both beam holes, and two surrounding regions. The X view regions are numbered from 1 to 5, and the Y view from 6 to 8. Each chamber is defined in this way to isolate the beam regions, which suffer from much higher rates. 146
- 67 The SOD means and widths in each DC region were fit to a Gaussian in a region close to the SOD peak in data and monte carlo. The above plot shows the ratio of fit SOD widths in the data to those in the monte carlo for the winter E799 run in $K_L \rightarrow 2\pi_D^0$. For a significant part of the winter run, the SOD widths were much higher in chamber 2 which, in addition to the other deviations, suggests a problem with the DC plane resolution modeling. 147
- 68 The SOD means and widths in each DC region were fit to a Gaussian in a region close to the SOD peak in data and monte carlo. The above plot shows the ratio of fit SOD means in the data to those in the monte carlo for the winter E799 run in $K_L \rightarrow 2\pi_D^0$. The SOD means were found to be lower by about 10 - 15 microns in the data. This suggests a systematic miscalibration of the XT maps. 148
- 69 The SOD means and widths in each DC region were fit to a Gaussian in a region close to the SOD peak in data and monte carlo. The above plot shows the ratio of fit SOD widths in the data to those in the monte carlo for the summer E799 run in $K_L \rightarrow 2\pi_D^0$. The deviations suggest a problem with the DC plane resolution modeling. 149
- 70 The SOD means and widths in each DC region were fit to a Gaussian in a region close to the SOD peak in data and monte carlo. The above plot shows the ratio of fit SOD means in the data to those in the monte carlo for the summer E799 run in $K_L \rightarrow 2\pi_D^0$. The SOD means were found to be lower by about 10 - 15 microns in the data. This suggests a systematic miscalibration of the XT maps. 150
- 71 The Z slope problem in the E799 winter $K_L \rightarrow 2\pi_D^0$ sample. The top plots show the vertex Z distribution with monte carlo overlay and the corresponding data/MC ratio with no special correction. The bottom plots show the same after a correction has been applied to the XT maps and the chamber resolutions. 151

- 72 The Z slope problem in the E799 summer $K_L \rightarrow 2\pi_D^0$ sample. The top plots show the vertex Z distribution with monte carlo overlay and the corresponding data/MC ratio with no special correction. The bottom plots show the same after a correction has been applied to the XT maps and the chamber resolutions. 152
- 73 The $K_L \rightarrow \pi^0 e^+ e^- \gamma$ decay final state. In $K_L \rightarrow \pi^0 e^+ e^- \gamma$ decay, two photons are the products of a π^0 decay while the third photon is from direct emission. The $e^+ e^-$ pair comes from an off-shell photon which converts internally. The decay dynamics are similar to those governing $K_L \rightarrow \pi^0 \gamma \gamma$ decay. 154
- 74 An example of $K_L \rightarrow 3\pi_D^0$ background to $K_L \rightarrow \pi^0 e^+ e^- \gamma$. In this case, a photon was lost from the Dalitz π^0 decay and a photon was lost from one of the normal π^0 decays. Since an intact π^0 survives, but the Dalitz π^0 decay is fragmentary, neither the cut on π^0 mass nor the cut against π^0 mass in the Dalitz case are effective here. Rather, cuts are made in the decay kinematics and in the RC counters to reduce this background. 155
- 75 $M_{ee\gamma}$, the invariant mass of the electron-positron pair and the lone photon, is plotted for $K_L \rightarrow 2\pi_D^0$ monte carlo and E799 data. The upper plot shows the $M_{ee\gamma}$ distribution for monte carlo events that satisfy the $K_L \rightarrow \pi^0 e^+ e^- \gamma$ analysis except for the cut against $M_{ee\gamma}$ in the best pairing hypothesis. While a cut at 145 MeV/ c^2 is sufficient on the high side, a cut must be placed at 115 MeV/ c^2 on the low side to eliminate $K_L \rightarrow 2\pi_D^0$ events with final state radiation, which comprise the low-side tail about 3.5 orders of magnitude down from the peak. The lower plot shows the same $M_{ee\gamma}$ distribution for E799 data satisfying the $K_L \rightarrow \pi^0 e^+ e^- \gamma$ analysis except for the cut against $M_{ee\gamma}$ in the best pairing hypothesis with $K_L \rightarrow 2\pi_D^0$ monte carlo overlaid. The normalization is absolute, and one can see the good agreement between data and monte carlo in the low side tail. 158
- 76 The algorithm used to identify π^0 decay products in this analysis is a simple one. There are three possible ways to make pairs of two photons given three photons. The pairing hypothesis with the invariant mass closest to the nominal π^0 mass is chosen as the best pairing. This calculation is done with 4-vectors assuming that the photons originate at the charged vertex. Occasionally, the “best” pairing hypothesis is not the correct pairing. In this case, the lone photon and $e^+ e^-$ pair may appear to have an invariant mass different from the π^0 mass. 159
- 77 The difference of $M_{ee\gamma}$ and M_{π^0} is plotted for $K_L \rightarrow 2\pi_D^0$ monte carlo and E799 data for the 2nd best pairing hypothesis. The upper plot shows the $|M_{ee\gamma} - M_{\pi^0}|$ distribution for monte carlo events that satisfy the $K_L \rightarrow \pi^0 e^+ e^- \gamma$ analysis except for the cut against $M_{ee\gamma}$ in the 2nd best pairing hypothesis. Since there is not as much of this background as in the best pairing case (see figure 75,) we can treat the $K_L \rightarrow 2\pi_D^0$ background peak as symmetric in the 2nd best pairing case. A cut at ± 10 MeV/ c^2 relative to the nominal π^0 mass is made to eliminate this mis-pairing background. The lower plot shows the same distribution for E799 data satisfying the $K_L \rightarrow \pi^0 e^+ e^- \gamma$ analysis except for the cut against $M_{ee\gamma}$ in the 2nd best pairing hypothesis with $K_L \rightarrow 2\pi_D^0$ monte carlo overlaid. 160

- 78 The difference of $M_{ee\gamma}$ and M_{π^0} is plotted for $K_L \rightarrow 2\pi_D^0$ monte carlo and E799 data. The upper plot shows the $|M_{ee\gamma} - M_{\pi^0}|$ distribution for monte carlo events that satisfy the $K_L \rightarrow \pi^0 e^+ e^- \gamma$ analysis except for the cut against $M_{ee\gamma}$ in the worst pairing hypothesis. Since there is not as much of this background as in the best pairing case (see figure 75,) we can treat the $K_L \rightarrow 2\pi_D^0$ background peak as symmetric in the 2nd best pairing case. A cut at $\pm 5 \text{ MeV}/c^2$ relative to the nominal π^0 mass is made to eliminate this mis-pairing background. The lower plot shows the same distribution for E799 data satisfying the $K_L \rightarrow \pi^0 e^+ e^- \gamma$ analysis except for the cut against $M_{ee\gamma}$ in the 2nd best pairing hypothesis with $K_L \rightarrow 2\pi_D^0$ monte carlo overlaid. 161
- 79 A table of simple ways $K_L \rightarrow 3\pi_D^0$ can contribute to $K_L \rightarrow \pi^0 e^+ e^- \gamma$ backgrounds. In addition to the cases listed above, more particles can be lost or fused and replaced by accidentals. However, such cases are more rare. 162
- 80 The distribution of CsI cluster shape chi-square for photons is plotted for $K_L \rightarrow 2\pi_D^0$ and $K_L \rightarrow 3\pi_D^0$ events. The upper plot shows the distribution of CsI cluster shape chi-square for $K_L \rightarrow 2\pi_D^0$ events that satisfy the $K_L \rightarrow \pi^0 e^+ e^- \gamma$ analysis except for the cut against $M_{ee\gamma}$ invariant mass in the best pairing hypothesis. Overlaid is a sample of $K_L \rightarrow 3\pi_D^0$ monte carlo events that satisfy the $K_L \rightarrow \pi^0 e^+ e^- \gamma$ analysis except for the cut against $K_L \rightarrow 3\pi_D^0$ kinematics and the final kaon mass cut. The excess in the case of $K_L \rightarrow 3\pi_D^0$ is due to the existence of overlapping photon clusters in the CsI. Since $K_L \rightarrow 3\pi_D^0$ with two extremely close fusions may pass mass and pt cuts, it is a serious background problem. A cut is made at a chi-square value of 5 to reject fusions. The lower plot shows the cluster shape chi-square variable for photons in a similarly defined $K_L \rightarrow 2\pi_D^0$ sample in the E799 data with $K_L \rightarrow 2\pi_D^0$ monte carlo overlay. The agreement between data and monte carlo is good. 164
- 81 Kinematic variable κ , defined in appendix D, for $K_L \rightarrow 3\pi_D^0$ rejection. The top plot shows E799 data in the $K_L \rightarrow \pi^0 e^+ e^- \gamma$ sample after all cuts (except this one and the cut on $e^+ e^- \gamma \gamma \gamma$ mass) with a $K_L \rightarrow 3\pi_D^0$ monte carlo overlay. The normalization is absolute, and as expected, the background at this stage is dominated by $K_L \rightarrow 3\pi_D^0$. The bottom plot shows the $K_L \rightarrow \pi^0 e^+ e^- \gamma$ monte carlo distribution in κ also with $K_L \rightarrow 3\pi_D^0$ monte carlo. The data at $\kappa = 0$ are populated by non-fragmentary $K_L \rightarrow \pi^0 e^+ e^- \gamma$ events. The cut effectively removes the threat of $K_L \rightarrow 3\pi_D^0$ background. . . 166
- 82 Kinematic variable κ , defined in appendix D, is plotted for E799 data with the cut against $K_L \rightarrow 2\pi_D^0$ removed with a $K_L \rightarrow 2\pi_D^0$ monte carlo overlay (top). The normalization is absolute, and the $K_L \rightarrow 2\pi_D^0$ shape is well predicted by the monte carlo. The bottom plot compares the distributions in κ for $K_L \rightarrow 2\pi_D^0$ and $K_L \rightarrow \pi^0 e^+ e^- \gamma$ monte carlo. 167

- 83 The distribution of maximum energy deposited in a spectrometer anti (SA) is plotted for $K_L \rightarrow 2\pi_D^0$ and $K_L \rightarrow 3\pi_D^0$ events. Energy is summed in each SA counter, and the maximum energy deposit over the set of SA counters is used to make these plots. The upper plot shows the distribution of maximum energy deposit in the SAs for $K_L \rightarrow 2\pi_D^0$ events that satisfy the $K_L \rightarrow \pi^0 e^+ e^- \gamma$ analysis except for the cut against $M_{ee\gamma}$ invariant mass in the best pairing hypothesis. Overlaid is a sample of $K_L \rightarrow 3\pi_D^0$ monte carlo events that satisfy the $K_L \rightarrow \pi^0 e^+ e^- \gamma$ analysis except for the cut against $K_L \rightarrow 3\pi_D^0$ kinematics and the final kaon mass cut. The $K_L \rightarrow 3\pi_D^0$ show a slight enhancement from lost photons striking an SA, while the $K_L \rightarrow 2\pi_D^0$ sample represents a null case for comparison. The lower plot shows a similarly defined set of $K_L \rightarrow 2\pi_D^0$ events in the E799 data with $K_L \rightarrow 2\pi_D^0$ monte carlo overlay. The agreement between data and monte carlo is excellent. However, the distinction between $K_L \rightarrow 3\pi_D^0$ and the null case is not so great. Therefore, we do not cut on the SAs. 168
- 84 The distribution of maximum energy deposited in a ring counter (RC) is plotted for $K_L \rightarrow 2\pi_D^0$ and $K_L \rightarrow 3\pi_D^0$ events. Energy is summed in each RC counter, and the maximum energy deposit over the set of RC counters is used to make these plots. The upper plot shows the distribution of maximum energy deposit in the RCs for $K_L \rightarrow 2\pi_D^0$ events that satisfy the $K_L \rightarrow \pi^0 e^+ e^- \gamma$ analysis except for the cut against $M_{ee\gamma}$ invariant mass in the best pairing hypothesis. Overlaid is a sample of $K_L \rightarrow 3\pi_D^0$ monte carlo events that satisfy the $K_L \rightarrow \pi^0 e^+ e^- \gamma$ analysis except for the cut against $K_L \rightarrow 3\pi_D^0$ kinematics and the final kaon mass cut. The $K_L \rightarrow 3\pi_D^0$ show a slight enhancement from lost photons striking an RC, while the $K_L \rightarrow 2\pi_D^0$ sample represents a null case for comparison. The lower plot shows a similarly defined set of $K_L \rightarrow 2\pi_D^0$ events in the E799 data with $K_L \rightarrow 2\pi_D^0$ monte carlo overlay. The agreement between data and monte carlo is excellent. A cut is at 0.5 MIPS (equivalent signal *Minimum Ionizing Particle*) to reject $K_L \rightarrow 3\pi_D^0$ background. 169
- 85 The E799 dataset is plotted with $K_L \rightarrow 3\pi_D^0$ and $K_L \rightarrow 2\pi_D^0$ monte carlo overlay after all cuts except for the cut against the kinematical variable, κ . The normalization is absolute. 170
- 86 The $K_L \rightarrow \pi^0 \gamma \gamma$ final state. This can cause a background problem for the $K_L \rightarrow \pi^0 e^+ e^- \gamma$ analysis if one of the direct photons converts to an $e^+ e^-$ pair in the vacuum window. This case is suppressed by the requirement that the decay vertex occurs upstream of the vacuum window. 171
- 87 An example of a double conversion event from $K_L \rightarrow 2\pi^0$ where two photons convert in the vacuum window. This can in principle contribute to $K_L \rightarrow \pi^0 e^+ e^- \gamma$ background. 172
- 88 A double conversion event in the charged spectrometer. In this scenario, two photons from a $K_L \rightarrow 2\pi^0$ or $K_L \rightarrow 3\pi^0$ event convert in the window. The opening angle between the tracks of each conversion pair is small, so that each pair looks like a single track in the upstream chambers. In the unlikely event that an oppositely charged pair of soft electrons is swept out by the analysis magnet, background can result since the resulting tracks may appear to make a good vertex upstream of the vacuum window. This background has been modeled and is very small, however, and does not contribute to the $K_L \rightarrow \pi^0 e^+ e^- \gamma$ sample at this level. 173

- 89 A natural four track event, such as $K_L \rightarrow 3\pi^0$ double Dalitz, has tracks that are well separated in the upstream chambers. Therefore, if two oppositely charged tracks get swept away at the analysis magnet, the background can be suppressed by making a cut against the possibility of in-time hit pairs in the upstream chambers. 174
- 90 For each event, we can find the minimum distance between the extrapolated position of an upstream segment of a charged track at the CsI and a hardware cluster in the CsI. This variable is plotted above for $K_L \rightarrow \pi^0 e^+ e^- \gamma$ monte carlo and for remaining events in the E799 data. The significance of this variable is that bremsstrahlung photons are distinguished by their closeness to the associated charged track. (Since charged tracks are bent by the analysis magnet, we use extrapolated upstream track segments for this calculation.) One can see a small peak of 9 events in the data that seem to correspond to bremsstrahlung events and are not explained by the $K_L \rightarrow \pi^0 e^+ e^- \gamma$ monte carlo. 175
- 91 The angle ϕ between the e^+e^- plane and the $\gamma\gamma$ plane is depicted above. To get the angle between the decay planes, the relation $\cos\phi = \frac{(\gamma_1 \times \gamma_2)(e_1 \times e_2)}{norm}$ 176
- 92 $\Delta\phi$ is plotted for two pairings of three photons. Seven of the nine events in the $D_{min} < 1.25$ cm region also satisfy $\Delta\phi < 0.2$ radians simultaneously in two of the angular variables (and in the third as well.) The top plot shows the distribution for $K_L \rightarrow \pi^0 e^+ e^- \gamma$ events and the bottom is the data. Therefore, the seven events are coplanar. This is consistent with $K_L \rightarrow e^+ e^- \gamma \gamma \gamma$ decay, aka $K_L \rightarrow e^+ e^- \gamma$ with two internal bremsstrahlung photons. 177
- 93 The $e^+e^- \gamma$ invariant mass distribution of the final $K_L \rightarrow \pi^0 e^+ e^- \gamma$ sample is plotted above (crosses) on top of a plot in signal plus data monte carlo (solid histogram). The shaded histogram shows the background only monte carlo. The normalization is absolute. 183
- 94 The $e^+e^- \gamma \gamma \gamma$ invariant mass distribution of the final $K_L \rightarrow \pi^0 e^+ e^- \gamma$ sample is plotted above (crosses) on top of a plot in signal plus data monte carlo (solid histogram). The shaded histogram shows the background only monte carlo. The normalization is absolute. 184
- 95 The shaded histogram shows the $\gamma\gamma$ invariant mass spectrum from the latest KTeV measurement of $K_L \rightarrow \pi^0 \gamma \gamma$ [Alavi 99b]. Overlaid with error bars is the $e^+e^- \gamma$ invariant mass spectrum from the $K_L \rightarrow \pi^0 e^+ e^- \gamma$ analysis. The agreement between the two spectra is very good. The cut on (non- π^0) $\gamma\gamma$ mass in the $K_L \rightarrow \pi^0 \gamma \gamma$ analysis is tighter than the corresponding cut on $e^+e^- \gamma$ mass in the $K_L \rightarrow \pi^0 e^+ e^- \gamma$ analysis. In the $K_L \rightarrow \pi^0 \gamma \gamma$ analysis, events with (non- π^0) $\gamma\gamma$ mass between 110 MeV/ c^2 and 160 MeV/ c^2 . $K_L \rightarrow \pi^0 \gamma \gamma$ plot courtesy of S. Taegar. 186
- 96 The ratio of χ^2 calculated using the approximate formula for the errors A.8 to that calculated using the general formula A.5. The ratios for the uncorrelated assumption are the circles, the ratios for the correlated assumption are the squares, and the ratios for the anti-correlated case are the triangles. Since the general formula A.5 is seen to give χ^2 results of about 1 per degree of freedom, a ratio of less than one indicates that the approximate formula is producing abnormally large errors. 194

- 97 A confidence belt constructed using the method of Feldman and Cousins in the μ - n plane. The belt constructed above is for the case of an exact background of 3.0 at a confidence level of 90%. The horizontal intervals bounded by the jagged curves are known as acceptance intervals; the n included in each acceptance interval satisfy relation B.1. A confidence interval at a specific n is the smallest interval in μ that contains all of the acceptance regions intersected by a vertical line drawn at n 198
- 98 The distributions of \mathcal{P}_1 versus m_2 for $K_L \rightarrow 3\pi_D^0$ and $K_L \rightarrow 2\pi_D^0$ Monte Carlo after requiring two charged tracks and 5 electro-magnetic clusters in the CsI calorimeter. The $K_L \rightarrow 3\pi_D^0$ events (top) are fragmentary and thus occupy an extended region in the \mathcal{P}_1 - m_2 plane. The $K_L \rightarrow 2\pi_D^0$ events are complete and fall along a kinematic boundary in the \mathcal{P}_1 - m_2 plane. A small fraction of $K_L \rightarrow 2\pi_D^0$ (bottom) events do in fact populate nether regions of the plot, but these are due to misreconstructed events or events with accidental energy. These poor events would fail to satisfy more stringent analysis cuts anyway. 206
- 99 The top plot shows the $K_L \rightarrow 3\pi_D^0$ and the $K_L \rightarrow 2\pi_D^0$ distributions in the $\log \mathcal{P}_1$ - m_2 plane. A line is drawn parallel to the cut in this plane that gives the best separation between the two distributions. Requiring in an analysis that the events be located above or below this line is equivalent to cutting on the variable $V_c \equiv \log \mathcal{P}_1 - 15m_2$ where the line in the plot corresponds to $V_c = 0$. One can plot the distribution of events transverse to this defining line. Such a plot is shown in the bottom plot where the same Monte Carlo distributions of $K_L \rightarrow 3\pi_D^0$ and $K_L \rightarrow 2\pi_D^0$ are overlaid. 207

*List of Figures

CHAPTER 1

Introduction to $K_L \rightarrow \pi^0 e^+ e^- \gamma$

The rare decay $K_L \rightarrow \pi^0 e^+ e^- \gamma$ is an interesting decay mode for two reasons. First, it occurs at a higher rate than its non-radiative partner, $K_L \rightarrow \pi^0 e^+ e^-$. The latter decay mode has been the object of many experimental searches recently since an observation of $K_L \rightarrow \pi^0 e^+ e^-$ may signal direct CP violation. In the kinematical region where the lone γ of $K_L \rightarrow \pi^0 e^+ e^- \gamma$ is soft, $K_L \rightarrow \pi^0 e^+ e^- \gamma$ may become a background problem for $K_L \rightarrow \pi^0 e^+ e^-$. The second interesting fact about $K_L \rightarrow \pi^0 e^+ e^- \gamma$ is that it shares a $\pi^0 \gamma \gamma$ intermediate state with the CP conserving contribution to $K_L \rightarrow \pi^0 e^+ e^-$. In the case of $K_L \rightarrow \pi^0 e^+ e^-$, both photons are off shell; while in the case of $K_L \rightarrow \pi^0 e^+ e^- \gamma$, only one is. Since there are three potential contributions to $K_L \rightarrow \pi^0 e^+ e^-$ including the CP conserving one, studies of CP violation using $K_L \rightarrow \pi^0 e^+ e^-$ depend upon a reliable calculation of the $K_L \rightarrow \pi^0 \gamma \gamma$ contribution. The calculation of $\text{BR}(K_L \rightarrow \pi^0 \gamma \gamma)$ has been carried out up to $O(p^6)$ in Chiral Perturbation Theory (ChPT), and the most recent results are in agreement with theory. A measurement of the branching ratio of $K_L \rightarrow \pi^0 e^+ e^- \gamma$ is an important check of these calculations. The analysis presented in this dissertation represents the first observation of the rare decay $K_L \rightarrow \pi^0 e^+ e^- \gamma$. The calculation of $\text{BR}(K_L \rightarrow \pi^0 e^+ e^- \gamma)$ and discussions of its importance appear in [Donoghue 97] and [Gabbiani 97].




Figure 1. The diagram representing CP conserving $\pi\gamma\gamma$ intermediate state to $K_L \rightarrow \pi^0 e^+ e^-$.

1.1 Introduction to $K_L \rightarrow \pi^0 e^+ e^- \gamma$

Curiously, the rate of $K_L \rightarrow \pi^0 e^+ e^- \gamma$ is higher than the rate of $K_L \rightarrow \pi^0 e^+ e^-$. While counterintuitive at first glance, the reason becomes obvious when the intermediate states are considered. Figure 1 shows the diagram corresponding to the CP conserving contribution to $K_L \rightarrow \pi^0 e^+ e^-$. If we compare this diagram with the leading diagram corresponding to the decay $K_L \rightarrow \pi^0 e^+ e^- \gamma$ in figure 2, we see that there is one less electro-magnetic coupling in the $K_L \rightarrow \pi^0 e^+ e^- \gamma$ decay, so that the amplitude should be on the order of the electro-magnetic coupling constant, α , larger than the $K_L \rightarrow \pi^0 e^+ e^-$ amplitude.

Before delving into the details of the calculation for the rate of $K_L \rightarrow \pi^0 e^+ e^- \gamma$, it is interesting to briefly review the physics behind $K_L \rightarrow \pi^0 e^+ e^-$ and its CP conserving $\pi^0 \gamma \gamma$ intermediate state.




Figure 2. Generalized diagram for the rare decay $K_L \rightarrow \pi^0 e^+ e^- \gamma$. If we compare this figure with figure 1, we see that the $K_L \rightarrow \pi^0 e^+ e^- \gamma$ process has one less internal electro-magnetic coupling. This means that the rate for $K_L \rightarrow \pi^0 e^+ e^- \gamma$ should be enhanced relative to the CP conserving contribution to $K_L \rightarrow \pi^0 e^+ e^-$ by a factor of the electro-magnetic coupling, α .

1.2 Why is $K_L \rightarrow \pi^0 e^+ e^-$ a CP -Violating Decay ?

The operator CP is the combination of the two discrete operators C and P . C stands for charge conjugation, and is the operator which takes a particle state into the anti-particle state. C reverses the sign of all additive quantum numbers and charges, but leaves spins and momenta unchanged. P stands for parity, and is the operator which reverses the sign of all spatial dimensions. P changes the sign of momenta, but leaves the spins and internal quantum numbers of a state unchanged. The discovery of separate C and P violations in the mid 1950s led to a revolution in the theory of the weak interactions [Marshak 93], and it was hoped that the combined operator CP would restore some of the lost symmetry to the theory. In 1964, however, Christenson, Cronin, Fitch, and Turlay found evidence [Christenson 64] for a small CP violation in the decays of long lived neutral kaons to 2π . Since that time, it has been an imperative of particle physics to explore the phenomenon of CP violation in order to discover its origin.

1.2.1 CP Symmetry and the Neutral Kaons

Neutral kaons are produced in eigenstates of definite strangeness by the strong interaction. The K^0 is assigned strangeness $+1$, so that the quark assignments are given by

$$|K^0\rangle = |\bar{s}d\rangle, \quad (1.1)$$

and

$$|\bar{K}^0\rangle = |\bar{d}s\rangle. \quad (1.2)$$

The neutral kaons are then not eigenstates of CP . Instead, we have

$$CP|K^0\rangle = |\bar{K}^0\rangle, \quad (1.3)$$

and

$$CP|\bar{K}^0\rangle = |K^0\rangle. \quad (1.4)$$

The CP operation on the K^0 and \bar{K}^0 is arbitrary up to a phase.

$$CP|K^0\rangle = e^{i\phi}|\bar{K}^0\rangle, \quad (1.5)$$

$$CP|\bar{K}^0\rangle = e^{-i\phi}|K^0\rangle. \quad (1.6)$$

This phase appears since the strangeness eigenstates of the neutral kaon are degenerate in the strong and electro-magnetic interactions, but are mixed by the weak interaction.¹ The phase convention used here is $\phi = 0$; although $\phi = \pi$ is also common.² If CP were a good symmetry of the weak interaction, then the eigenstates of the weak interaction would also be CP eigenstates. With our phase convention, the CP eigenstates are given by the simple linear combinations

$$|K_1\rangle = \frac{1}{\sqrt{2}}(|K^0\rangle + |\bar{K}^0\rangle), \quad (1.7)$$

¹The phase that appears in 1.6 is proportional to the strangeness, so that the K^0 and \bar{K}^0 get re-phased oppositely.

²For a good discussion of all of the arbitrary phases and how they are related to measurable quantities, see [D'Ambrosio 96] or [Sachs 87] and references therein.

$$|K_2\rangle = \frac{1}{\sqrt{2}}(|K^0\rangle - |\bar{K}^0\rangle), \quad (1.8)$$

where the K_1 is the $CP = +1$ eigenstate and the K_2 is the $CP = -1$ eigenstate. The definite CP of the proposed decay states has an important consequence when calculating the lifetime of each state. If CP were a good symmetry, the K_2 could never decay to 2π , and the K_1 could never decay³ to 3π . The lower Q value of the 3π decay restricts the overall amount of phase space available for the K_2 decay; therefore, the K_2 has a longer lifetime than the K_1 [Gell-Mann 55].⁴ This prediction of a long-lived neutral kaon state was soon confirmed in [Lande 56]. The observed long-lived neutral kaon lifetime, τ_L , is $5.17 \pm 0.04 \times 10^{-8}$ seconds, while the short-lived neutral kaon lifetime, τ_S , is $0.8934 \pm 0.0008 \times 10^{-10}$ seconds [PDG 98].

As a result of the small violation of CP symmetry observed in $K_L \rightarrow 2\pi$ decay, it is no longer valid to assume that the CP eigenstates are the same as the weak interaction eigenstates as we did above in 1.7 and 1.8. However, unlike the maximality of individual P and C violation in the weak interaction, CP violation is a small effect. Therefore a reasonable ansatz is that the weak eigenstates are small perturbations on the exact CP states given in 1.7 and 1.8.

$$|K_L\rangle = \frac{1}{1 + |\epsilon|^2}(|K_2\rangle + \epsilon|K_1\rangle), \quad (1.9)$$

$$|K_S\rangle = \frac{1}{1 + |\epsilon|^2}(|K_1\rangle + \epsilon|K_2\rangle), \quad (1.10)$$

where $|K_L\rangle$ now represents the long lived neutral kaon and $|K_S\rangle$ is the short lived one and ϵ is a complex number parameterizing the admixture of states. Table 1 lists the different ways of thinking about neutral kaons discussed so far. Note that the states of constant lifetime are no longer orthogonal in this representation, but rather $\langle K_S | K_L \rangle = 2\text{Re}(\epsilon)$.

³The 2π state is always CP even because of the Bose symmetry of the system. The 3π state is CP odd when the total orbital angular momentum of all the parts is 0. This is the case in K decay because of energy considerations and the angular momentum barrier.

⁴Gell-Mann and Pais predicted the long lived neutral kaon on the basis of C invariance alone. However, the formalism is exactly the same if CP is assumed instead.

Basis	Definition	Interesting Trivia
K^0, \bar{K}^0	Well Defined Strangeness	Strong Interaction Production States
K_1, K_2	Well Defined CP	Conceptually Useful Approximate States
K_L, K_S	Weak Eigenstate	Definite Mass and Decay Lifetime

Table 1. A Table showing different ways of thinking about the neutral kaon.

There are two categories of CP violation. *Indirect* CP violation is distinguished by the above admixture of states and is parameterized by ϵ . *Direct* CP violation is a direct decay of one definite CP state into a state of opposite CP . For example, the decay $K_2 \rightarrow 2\pi$ would be direct CP violation. Until recently, there has been no clear evidence of direct CP violation.⁵

Thus, loosely speaking, there are two reasons why $K_L \rightarrow \pi^0 e^+ e^-$ is a CP -violating decay, and one reason why it is not. There is a direct CP violating contribution, an indirect CP violating contribution, and a CP conserving one.

1.2.2 CP Conserving Contribution to $K_L \rightarrow \pi^0 e^+ e^-$

The richest of the three contributions to $\text{BR}(K_L \rightarrow \pi^0 e^+ e^-)$ is the CP conserving one. This is because the physics involved highlights a ChPT calculation enhanced by vector meson exchange and also because the calculation of this contribution is related to predictions for two other decay modes, $K_L \rightarrow \pi^0 \gamma \gamma$ and $K_L \rightarrow \pi^0 e^+ e^- \gamma$. Because of this fact, we will introduce the discussion with a brief summary of $K_L \rightarrow \pi^0 \gamma \gamma$.

Some early attempts to calculate the rate of $K_L \rightarrow \pi^0 \gamma \gamma$ were done at $O(p^4)$ in chiral perturbation theory. While the $O(p^4)$ ChPT calculations correctly modeled the $\gamma \gamma$

⁵Recently, the KTeV collaboration has observed the first clear evidence of direct CP violation in the $K \rightarrow 2\pi$ system with a non-zero measurement of the parameter $\frac{\epsilon'}{\epsilon}$.

invariant mass ($m_{\gamma\gamma}$) spectrum, the measured branching ratio was more than a factor of 2 higher [Barr 2],[Papadimitriou 91], [Ecker 87]. Other models supplemented the ChPT calculation with vector meson exchanges [Heiliger 93], such as from ρ or ω . These models are known as vector meson dominance (VMD) models. The VMD models were marked by their higher overall predictions for the rate $K_L \rightarrow \pi^0 \gamma\gamma$, although the spectra of $\gamma\gamma$ invariant mass were wildly different from pure ChPT. Figure 5 shows diagrams representing the VMD contributions. Ecker pointed out that the VMD exchanges could be incorporated into the ChPT calculation at higher orders in the chiral expansion [Ecker 90]. At $O(P^6)$, the VMD effects are parameterized with a constant, a_v . Surprisingly, this approach is able to reproduce the observed rate *and* maintain the distinctive $m_{\gamma\gamma}$ spectrum [Ecker 93], [D'Ambrosio 97]. The recent KTeV measurement of the rate is $(1.69 \pm 0.07 \pm 0.08) \times 10^{-6}$ with $a_v = -0.78 \pm 0.06$ [Alavi 99b]. The calculated spectrum matches the data very well, so that we have confidence in applying the results to $K_L \rightarrow \pi^0 e^+ e^-$. We will use the latest model with the new value of a_v for estimating the CP conserving rate for $K_L \rightarrow \pi^0 e^+ e^-$.

The CP conserving contribution to $K_L \rightarrow \pi^0 e^+ e^-$ proceeds through the $\pi^0 \gamma\gamma$ intermediate state. A generalized diagram of this contribution is shown in figure 1. At lowest order in momentum, the photons form a spin-0 state. In this case, the coupling of the photons to $e^+ e^-$ is suppressed by a factor of m_e^2 due to the helicity conserving nature of the electro-magnetic vertices. At higher powers of momentum this is no longer true, but there is an angular momentum barrier to overcome to get those higher powers of momentum. These considerations led to the early belief that the CP conserving contribution to the $K_L \rightarrow \pi^0 e^+ e^-$ rate would be on the order of 10^{-13} [Donoghue 87], using a pure ChPT calculation at $O(p^4)$. However, as pointed out above, the CP -conserving 2 photon intermediate state can be enhanced via VMD effects. The VMD contributions are not helicity suppressed [Isidori 99], [Heiliger 93]. The VMD contributions are depicted in figure 5. The dependence of the CP conserving contribution to $K_L \rightarrow \pi^0 e^+ e^-$ on a_v is shown in figure 3, reproduced from [Gabbiani 97] and [Donoghue 95]. Using the value measured in KTeV, a value of about 2.0×10^{-12} is indicated on the graph.

$$BR(K_L \rightarrow \pi^0 e^+ e^-)_{CP \text{ Conserving}} \simeq 2.0 \times 10^{-12}, \quad (1.11)$$




Figure 3. The dependence of BR on a_v . *Reproduced courtesy of F. Gabbiani.*

using the central value for a_v . A fit done to existing data in [Gabbiani 97] favors a value of -0.96 for a_v , corresponding to an estimate of 4.9×10^{-12} for the CP conserving contribution to $K_L \rightarrow \pi^0 e^+ e^-$.

Both of the CP violating contributions to $K_L \rightarrow \pi^0 e^+ e^-$ are also small on the order of 10^{-11} - 10^{-12} in the branching ratio [Donoghue 95].




Figure 4. Diagram detailing the $K \rightarrow 3\pi$ contribution in chiral perturbation theory to $K_L \rightarrow \pi^0 e^+ e^- \gamma$. (For $K_L \rightarrow \pi^0 e^+ e^- \gamma$, one of the final photons is off shell.)

1.3 $K_L \rightarrow \pi^0 e^+ e^- \gamma$ and Its Relationship to $K_L \rightarrow \pi^0 e^+ e^-$

Interestingly, the rare decay $K_L \rightarrow \pi^0 e^+ e^- \gamma$ occurs at a higher rate than the decay $K_L \rightarrow \pi^0 e^+ e^-$. While both decays occur through a two photon intermediate state, the $K_L \rightarrow \pi^0 e^+ e^-$ decay must proceed with one more power of the electro-magnetic coupling constant, α , than the corresponding $K_L \rightarrow \pi^0 e^+ e^- \gamma$ decay⁶ [Donoghue 97]. Furthermore, since one of the photons is off-shell in the internal state of $K_L \rightarrow \pi^0 e^+ e^- \gamma$, the chiral calculations come in with different relative weights as in the calculation for $K_L \rightarrow \pi^0 \gamma \gamma$. Thus, a measurement of the branching ratio for $K_L \rightarrow \pi^0 e^+ e^- \gamma$ is a check on the new $O(p^6)$ calculation done for that decay mode [Donoghue 97], [Gabbiani 97].

The calculation for $K_L \rightarrow \pi^0 e^+ e^- \gamma$ was presented at two orders of chiral perturbation theory, $O(p^4)$ and $O(p^6)$, so that a comparison can be made. Some contributing diagrams are given in figure 4 for chiral perturbation theory and in figure 5 for vector meson exchange contributions. The vector meson exchange contributions, appearing at $O(p^6)$ of the calculation, enhance the overall decay rate. The prediction for the branching ratio is

⁶Although this point seems obvious, it is worthwhile to point this out. The topology that usually comes to mind involves an inner-bremsstrahlung photon and one *more* power of α .




Figure 5. Two diagrams detailing vector meson exchange contributions to $K_L \rightarrow \pi^0 e^+ e^- \gamma$. These contributions only appear at $O(p^6)$ in the calculation. (For $K_L \rightarrow \pi^0 e^+ e^- \gamma$, one of the final photons is off shell.)

$$BR(K_L \rightarrow \pi^0 e^+ e^- \gamma)_{O(p^6)} = 2.4 \times 10^{-8} \quad (1.12)$$

at $O(p^6)$ of chiral perturbation theory. At $O(p^4)$, the rate is given by

$$BR(K_L \rightarrow \pi^0 e^+ e^- \gamma)_{O(p^4)} = 1.0 \times 10^{-8} \quad (1.13)$$

The error in this calculation can be estimated from the errors in the basic coupling constants a_1 and a_2 ([Gabbiani 97], [Kambor 91]) for the decay $K_L \rightarrow 3\pi$. This gives an error in the above calculations of 10% [Gabbiani 99].

A prediction is also made for the $e^+ e^- \gamma$ invariant mass spectrum, which should look a lot like the $\gamma\gamma$ spectrum in $K_L \rightarrow \pi^0 \gamma\gamma$.

1.4 Experimental Technique

The basic analysis strategy involves the complete reconstruction of kaon decays : charged particle momentum and trajectory reconstruction provided by the charged particle spectrometer, electron-photon energy reconstruction by the cesium-iodide (CsI) calorimeter, electron identification by the CsI calorimeter and the transition radiation detector (TRD), and hermetic detection of escaped decay products by the veto system.

KTeV also provides for muon reconstruction. However, the muon systems are not used in this analysis and are described elsewhere, [KTeV 113].

The signal $K_L \rightarrow \pi^0 e^+ e^- \gamma$ events will be normalized to the decay $K_L \rightarrow 2\pi_D^0$ for the branching ratio measurement.⁷ $K_L \rightarrow 2\pi_D^0$ has the same final state as $K_L \rightarrow \pi^0 e^+ e^- \gamma$ but is distinguished mainly by the fact that the $e^+ e^- \gamma$ will make a π^0 mass. The principle backgrounds come from the normalization mode, $K_L \rightarrow 2\pi_D^0$, and from the more copious $K_L \rightarrow 3\pi_D^0$ events which are missing at least two photons. In order to remove the first background, we will rely mainly on mass cuts to remove those events where the $e^+ e^- \gamma$ invariant mass is close to the π^0 mass. In order to remove the latter background, which is fragmentary, we will rely on event quality cuts and a final kinematic cut.

⁷The π^0 decays to 2γ about 98.8% of the time. About 1.2% of the time, the π^0 decays to $e^+ e^- \gamma$. This is known as the Dalitz channel and is explicitly represented in this thesis by π_D^0 . In this vein, $K_L \rightarrow 2\pi_D^0$ is shorthand for $K_L \rightarrow 2\pi^0, \pi^0 \rightarrow e^+ e^- \gamma$.

CHAPTER 2

Introduction to the Beam and Detector

KTeV was designed to accept a large flux of neutral kaons for studies of direct CP violation through measurement of $\frac{\epsilon'}{\epsilon}$ in the $K \rightarrow 2\pi$ system (Fermilab Experiment 832) and through direct searches for CP violating rare decays of the K_L , such as $K_L \rightarrow \pi^0 e^+ e^-$ (Fermilab Experiment 799). The KTeV detector hall was located along the neutrino-muon (NM) beam-line in the fixed target area of Fermilab. The geography of KTeV hall with respect to other Fermilab landmarks is shown in figure 6. The E799 data were taken during two periods in 1997. The first running period started in January, 1997 and lasted until March, 1997 and is called the E799 “winter” run. The second running period started in July, 1997 and ran until early September, 1997 and is called the “summer” run.

In this chapter, I will describe the production and characteristics of the kaon beam and the hardware employed to observe kaon decays.

2.1 The KTeV Beam

Figure 7 shows the layout of the KTeV secondary beam system. The primary proton beam delivered by the Tevatron strikes a BeO target defined to be at the origin of the KTeV coordinate system. Neutral kaons are produced along with many other species of charged and neutral particles. The beams are refined by a series of sweeping magnets and collimators until an optimal kaon beam is produced. All of the sweeping magnets have fields oriented along the vertical, so that charged particle tracks bend in a horizontal




Figure 6. A sketch of the fixed target area at Fermi National Accelerator Laboratory. The KTeV experiment is located on the neutrino-muon (NM) beam-line. Note that the map given in the Fermilab telephone directory is wrong. Not shown : Rosati's.

plane. Design details of the KTeV secondary beam can be found in references [KTeV 97] and [KTeV 213].

2.1.0.1 The KTeV Coordinate System

The official KTeV coordinate system is defined so the the positive Z axis lies along the vector from the center of the target to the center of the upstream face of the CsI calorimeter. The positive Y axis is taken to lie along the vertical. The positive X axis is related to the other two by the right hand rule, so that looking from upstream to downstream and standing on your head, the positive X axis is on your right side.

2.1.1 The KTeV Secondary Beam

The KTeV beam starts deep within a steel cave containing the target pile. 800 GeV/c protons from the Fermilab Tevatron are delivered onto a target consisting of about 1.1 interaction lengths of BeO. The target was 3 mm square facing the beam, and 30 cm long. The primary proton beam was monitored by several beam shape monitors which helped the operators focus the beam onto our target. The closest such monitor was just upstream of our target (NM2SEED2). The beam spot had a 1 mm diameter full width, and the target was cooled by a moderate flow of nitrogen gas to maintain an operating temperature of about 125° C. During normal operation, the Tevatron delivered between




Figure 7. KTeV beam elements. Kaons are produced by 800 GeV/c protons on a BeO target at $Z=0$ m. Charged particles are removed from the beam at various points by sweeping magnets while the neutral beam is shaped by a set of collimators. A lead absorber in the beam-line serves to convert photons in the beam to e^+e^- pairs, which are later swept out. The common absorber was not used in E799.

2.5×10^{12} and 5.0×10^{12} protons onto the target. Immediately downstream of the target pile, beam cleanup begins. The remaining primary proton beam is deflected into a water cooled copper beam dump by the “target sweeper” (NM2S1) magnet at $Z=2$ m. The BeO target, target sweeper magnet, and beam dump all reside in the steel cave and are collectively known as the target pile. With respect to the axis of the KTeV detector, the primary beam nominally makes an angle of 4.8 mrad in the vertical plane in order to avoid the peak neutron fluxes at 0.0 mrad and achieve the best kaon to neutron ratio with high kaon flux.

Downstream of the target pile are three more sweeper magnets to remove charged particles from decays and from interactions. As the neutral kaon beam is refined in this way, it is also collimated by a series of more and more precise collimators. Charged particles are swept away by a magnet at $Z=14$ m from the target, the “E8 Hyperon” (NM2S2) magnet. The primary collimator, followed this magnet at $Z=21$ m. It is a steel and brass mask roughly defining the two square KTeV beams¹ separated from each

¹The were two beams because of the requirements of E832. In E832, one of the beams strikes a regenerator to make K_S . In 799, neither beam had a regenerator, and the beams were considered identical.

Beam Element	Z position (m) of Upstream Face	Transverse Size (m) or Aperture Size (m)
BeO Target	0.0 (at center)	0.003 square (4.8 mRad angle with vertical)
Target Sweeper	2.0	(Not Limiting)
“E-8 Hyperon” magnet	14.0	(Not Limiting)
Pb Absorber	18.5	(in beam)
Primary Collimator	20.0	1.6 mrad beam separation
“ μ sweep 2” magnet	22	(not limiting)
Vacuum Begins	50	
Slab Collimator	40.0	(winter only)
Defining Collimator	85.0	4.4 cm winter square beams
	85.0	5.2 cm summer square beams
Final Sweeper	90.0	(fringe fields up to 94.0 m)

Table 2. Z positions and aperture information for important elements in the KTeV secondary beam system. Defining apertures are listed in **bold**.

other by 1.6 mrad. In front of the primary collimator at $Z=18.5$ m, a 7.5 cm long Pb absorber sits in the neutral beams in order to convert photons in the beam to charged e^+e^- pairs. Another magnet at $Z=22$ m, the “ μ -sweep 2” magnet” (NM2S3), removed charged particles from the beam. Just downstream of the primary collimator, at about $Z=22$ m, the particle beams enter a vacuum region which extends all the way to the end of the decay region, $Z=158.9$ m.

Jumping ahead to $Z=85.0$ m, there is a tungsten defining collimator which sharply defines the final shape of the neutral beams. For E799, each beam was 4.4 cm square at the defining collimator during the winter run and was 5.2 cm square at the defining collimator during the summer run. (To achieve the change, the winter defining and

primary collimators were swapped with new ones for the summer.) The defining collimator was designed to be as close to the main K_L decay volume as possible. At $Z=40.0$ m, between the primary and defining collimators, there was a 2m long steel slab collimator located between the two beams. This collimator made it impossible for scattered particles from one beam to cross over to the other beam. The slab collimator was in place during the winter run, but was removed for the summer run when larger beams were used. When the experiment was down for repairs, a beam stop was rolled in at $Z=50$ m. The beam stop was required to be in place before any personnel could be granted access to the detector hall.

Downstream of the defining collimator, at $Z=90.0$ m, there was a “final sweeper” magnet that removed any surviving charged particles from interactions and upstream decays. Inside of this magnet, the single beam pipe split into two 4 inch square beam pipes for transport to the main KTeV decay volume. The final sweeper was about 3 meters long, but fringing fields extended out to about 94 m.

There was also a spin-rotator magnet (important for several hyperon analyses,) and jaw collimators upstream of the beam stop. These components are not germane to this analysis. Table 2 shows a listing of the Z positions of the various elements in the KTeV beam.

2.1.2 The Timing of the KTeV Beam

The K_L beam didn’t arrive continuously. During the KTeV running, our minutes were divided approximately into 20-23 second “on-spill” periods followed by 38 second “off-spill” periods. During a typical off-spill period, proton current inside of the Tevatron was ramped up to the required energy and intensity. During a typical on-spill period, a portion of the Tevatron current was extracted onto our target. The beam during the on-spill period had micro-structure to it as well. The spill was divided into 19 nsec buckets. During each bucket, the beam arrived in a 1-2 ns pulse. A 53 MHz signal from the accelerator enabled us to synchronize the trigger with respect to the buckets. Lastly, the neighboring NUTeV neutrino experiment required high intensity “pings” of

short duration. Delivered at the beginning of the on-spill period, the KTeV detector was gated off during the pings which were associated with high rates in our detectors.

During the spill, it might happen that the proton intensity varied from bucket to bucket. This effect can lead to non-uniform accidental rates and therefore a possible non-uniform acceptance from bucket to bucket. As a measure of the bucket non-uniformity during the spill, the bucket occupancy autocorrelation (BOA) was monitored.

$$\text{BOA} = \frac{\langle I_b \cdot I_b \rangle}{\langle I_b \rangle^2}, \quad (2.1)$$

where I_b is the intensity in a given bucket b . Non-uniformity is signaled by a deviation of the BOA above 1.0.

The BOA was estimated by the ratio R of the number of coincidences between two uncorrelated detector elements, $D1$ and $D2$, during a given spill and the number of coincidences obtained when one of the sources is delays by several buckets.

$$R = \frac{\sum_b D1_b \cdot D2_b}{\sum_b D1_b \cdot D2_{(b+\Delta)}}, \quad (2.2)$$

where b denotes a given bucket and Δ is a suitable delay of several buckets. If the event rates between the in-time bucket and the delayed bucket are uncorrelated, then R is a measure of the *BOA*. The detector elements (described below) used in the determination of the BOA were the 90° target counter and the V trigger hodoscope plane. The BOA was monitored in the KTeV control room as well as the beam operators' control room. If the BOA deviated significantly from 1.0, the beam operators made adjustments to bring the beam back into line. During the KTeV run, the BOA was typically between 1.1 and 1.6 which was considered acceptable. If the BOA was > 2.0 , we complained to the nice people in the main control room.

2.1.3 Final Characteristics of the KTeV Beam

The action of the sweeping magnets deflected charged particles out of the beams and away from the detector. The KTeV detector therefore sat between two “lobes” of charged particles, mainly muons. The magnets were able to deflect muons away from the detector




Figure 8. Simulated K_L momentum in E799. The top plot shows the raw generated kaon momentum spectrum derived from the Malensek spectrum. The bottom plot shows the K_L momentum spectrum for K_L decays which satisfy the triggering requirements for the analysis considered in this thesis. Decays of K_L particles below about 20 GeV/c fail to satisfy triggering requirements because they tend to decay upstream and have decay products that escape the detector.

up to 600 GeV/c. The remaining muons were either above 600 GeV/c and very rare or were lower momentum muons that scattered back into the detector.

The production of neutral kaons is described by the Malensek spectrum, described in chapter 6. Using this parameterization, the mix of K^0 and \bar{K}^0 was approximately 54% K^0 to 46% \bar{K}^0 . The momentum spectrum of K_L is shown in figure 8.

The canonical KTeV decay region follows the final sweeper immediately. Upon entering the main decay volume, the beams of E799 contained K_L particles with neutrons, and small amounts of cascades, lambdas, and K_S . Although many hyperon lifetimes away

Particle	Flux Relative to K_L
n	3.0 - 3.5
K_L	1.0
Λ	0.02
Ξ	0.0008

Table 3. Relative populations of various particle species surviving in the E799 beam.

from the target, KTeV has managed to do much interesting physics with the lambda and cascade survivors. In particular, the cascade beta decay has been observed for the first time by KTeV. The neutrons were more numerous than the K_L . The ratio of neutrons to K_L particles was about 3:1 during E799. (This ratio was different in E832 because of the addition of the common absorber to the beam-line.) The other neutral particles, such as lambda or cascade hyperons, are present in much smaller amounts, suppressed by several orders of magnitude. Table 3 shows the estimated relative populations of various particle species surviving in the E799 neutral beams [KTeV 141]. The K_S particles which decay in our volume which survive from the target are all at very high momentum, typically 200 GeV/c. During optimal running, E799 got 0.5 - 1.0 K_L particles per bucket.

2.2 The KTeV Detector

The core of the KTeV detector, located in the KTeV detector hall at NM4, is a charged particle spectrometer followed by an electro-magnetic calorimeter and triggered by a pair of thin scintillator hodoscopes. Other detectors include the photon vetoes necessary to define a hermetic region of acceptance, and transition radiation detectors (TRDs) to provide electron identification. However, the K_L particles in which we are interested decay in the vacuum decay tank, so we will begin our survey of the KTeV detector there.




Figure 9. The KTeV detector. The scale in the Z direction is compressed with respect to the X direction by a factor of about 11.

2.2.1 The KTeV Vacuum Decay Tank and Ring Counter Vetoes

Upon exiting the final sweeper magnet, the two neutral beams enter the vacuum decay tank. The decay tank is an airtight steel enclosure that begins² at the upstream end at $Z=93.0$ m and continues downstream to 158.9 m. At the upstream end, the tank is about 15 cm in diameter; while at the downstream end, the tank is about 1.5 m in diameter. Thus, the vacuum decay tank resembles a giant refracting telescope. The pressure inside the decay tank was kept at 10^{-6} torr during the running of E799. At the downstream end of the decay volume, air was prevented from rushing into the vacuum decay tank by a strong, thin mylar-kevlar window. The thickness of this window was measured to be approximately 0.0015 radiation lengths. Just downstream of the vacuum window, a space was provided for the lowering of a safety shutter which was required to be down whenever personnel were in the detector hall.

²Note that the actual vacuum begins just downstream of the primary collimator. The beam is conveyed by a series of evacuated pipes to the main decay tank.

Detector Element	Z position (m) at Upstream Face	Relevant Aperture Size (m)
Defining Collimator (winter)	85.0	0.0425 square each (0.06855 cm separation)
Defining Collimator (summer)	85.0	0.0503 square each (0.06855 cm separation)
Final Sweeper	90.0	(Not Limiting)
E799 Main Decay Volume Begin	93.0	0.457 diameter
RC6	132.6	0.84 square
RC7	138.6	0.84 square
RC8	146.6	1.18 square
RC9	152.6	1.18 square
RC10	158.6	1.18 square
Vacuum Window	158.89	96 in. diameter

Table 4. Z positions and relevant apertures for elements in and around the KTeV main decay volume. The limiting apertures are in **bold**.

Inside the vacuum of the main decay tank were a series of photon vetoes called the ring counters (RCs) [KTeV 124]. The RC inner apertures were square. Each RC was designed to detect photons escaping from the detector at energies down to 100 MeV while providing hermetic coverage of escaping particles in the decay volume. The design of each RC calls for 16 radiation lengths of Pb-scintillator sandwich to achieve this. The first 16 layers consisted of a sheet of 0.5 radiation length Pb and scintillator, while the remaining 8 layers had 1.0 radiation lengths of Pb and scintillator. Light was guided out of the RC modules by glass fibers to phototubes mounted around the outer edge of each counter. The phototube signals were digitized by LeCroy 4300 ADCs. Discriminated signals from the RCs also provided source to the Level 1 trigger to quickly veto an event on the detection of decay products that would otherwise escape the detector. These

sources were set to fire if any RC registered more than 0.5 GeV of energy and provided this information on a bucket by bucket basis. Table 4 shows the Z positions of the various elements in the decay volume relative to target with the aperture information.

2.2.2 The Drift Chamber Spectrometer System

Located just downstream of the vacuum window is the charged particle spectrometer. The purpose of the spectrometer is to find the trajectories of charged particles and to measure the momentum of each charged track. By tracing the trajectories upstream to a point of closest approach, the position of the decay vertex may be inferred. In addition to these responsibilities, fast signals from the upstream chamber are used to provide input to the Level 1 trigger, and hit counting logic is implemented to help find events that are likely to reconstruct tracks in the Level 2 trigger system.

The KTeV spectrometer consists of four drift chambers split by an analysis magnet. Lots of experience exists with this set of drift chambers, which were first used in Fermilab E731 and later in E773 and E799-I, [Patterson 92], [Gibbons 93a], [Briere 95]. (The main difference is that aluminum based wires are used in KTeV to reduce scattering.) The spaces between the drift chambers were taken by plastic bags filled with helium. The purpose of the helium bags was to reduce multiple scattering of charged tracks during the momentum measurement. It was determined after the run, however, that a significant amount of air had leaked into the helium bags during the run, [Zimmerman 99].

2.2.2.1 *The Drift Chambers*

There were four drift chambers in the KTeV spectrometer numbered DC1 through DC4 from upstream to downstream. DC1 was the smallest chamber in cross section presented to the beam, while DC4 was the largest. Figure 10 shows the inner geometry of each drift chamber. Each chamber has four sense planes for the detection of amplified signals from passing charged particles. The first pair is oriented along the y direction (“X view”) in the upstream section of each chamber while the second is oriented along




Figure 10. The drift chamber geometry looking along the positive Y axis. The open boxes represent the field shaping wires and the crosses represent the sense wires. (The solid lines represent field shaping wires and the dashed lines represent the sense plane in the rotated view.) The upstream pair of sense planes consists of wires oriented along the Y direction to provide a measurement of X position. The downstream pair of sense planes consists of the planes along the X direction for a measurement of Y position. Adjacent sense wires are 12.7 mm apart. The two sense planes in each view are staggered to allow for unambiguous position measurement.

the x direction (“Y view”). The sense wires were 1 mil gold-plated tungsten and the field wires were 4 mil gold-plated aluminum. The field shaping wires were arranged in a hexagonal pattern around each sense wire, making an approximately cylindrical drift cell. The window guard wires were also 4 mil gold-plated aluminum, and were run at the same voltage as the field shaping wires to maintain electric potential at the edges of adjacent drift cells. Each sense plane in a given view and chamber is staggered by 6.35 mm with respect to its partner (one-half cell) to allow for unambiguous position measurement.

The position measurement starts when a charged track deposits ion pairs in the ar-

gon/ethane gas mixture. The ions then drift at an approximately constant velocity until they reach the singular field around a sense wire at which point they are amplified in the gas into a measurable signal. The drift velocity in the KTeV drift chambers was approximately 5 cm/usec at the normal operating voltage of -2450 V to -2600 V on the cathodes and window wires. The voltages were chosen to achieve the stated drift velocity and the necessary gain to see a signal with 99% efficiency.³ The drift chambers were run with a 49.75% argon - 49.75% ethane gas mixture with 0.5% isopropyl alcohol by volume used to increase the useful life of the chambers. In the summer run, the proportion of alcohol was raised to 1.0%. Alcohol in the chamber gas mix helps absorb damaging ultraviolet light produced during gas amplification, helping to prevent chamber aging.

The signals were amplified and discriminated by chamber mounted preamplifier cards at fixed gain. The output of the discriminators was fed into channels of LeCroy 3377 TDC (time to digital converter) operated in common-stop mode, so that signals reaching a sense wire start the TDC counting in time. All of the running TDCs are then stopped by a common signal from the Level 1 trigger system. A typical distribution of such hit times appears in figure 11. The window identified with in-time activity in the drift chambers is taken to be the region $115\text{ns} < t < 350\text{ns}$. The abrupt turn-on at 350 ns in the plot corresponds to short drift times near the sense wire while the tail at 200 ns corresponds to longer drift times far from the sense wire. The drift times are converted to distances using the “XT” maps. These are the conversion relations involved in turning drift times into distances and are calibrated throughout the run. If there were only one drift plane per view, however, there is an ambiguity in that one does not know whether the hit lies above or below the hit wire. The second staggered drift plane in each view gives us another measurement that allows us to resolve the ambiguity. This is illustrated in figure 10.

A further requirement can be made on hit pairs. For a single track and pair of hits along that track, the sum of distances, or *sod*, for that hit pair must add up to a half cell

³The drift time actually varies across a drift cell, and it is most pronounced in low field regions lying between any two adjacent field wires. The geometry of the wire chambers tries to minimize this effect. The variation in drift times across a cell can be calibrated to produce a map of distances corresponding to measured TDC times.




Figure 11. The arrival times of signals in DC2 during run 8397. Operated in common-stop mode, the hits close to a wire start the TDC counting earlier than other hits and contribute to the “bluff” near 350 ns. Later hits trail off to the left.

width, or 6.35 mm. Figure 12 shows the distribution of sods from the first drift chamber X-view for identified $K_L \rightarrow 2\pi_D^0$ normalization mode events. The position resolution of the drift chambers can be obtained from the sod distribution. The sod distribution has a Gaussian width of about 150 microns. Since two hits are used to make a sod measurement, the individual position resolution of the drift chambers is given by $\frac{150}{\sqrt{2}}\mu\text{m} \simeq 100\mu\text{m}$. The exact position resolution of each plane varied during the run as a function of applied voltage and gas composition, and was monitored throughout the run.

During the running of KTeV, an anomaly was observed in the sod distributions in the form of a high-sod tail. The appearance of the high-sod tail is position dependent and may be rate dependent. Maps of high-sod probability were made for E832. While that work is currently in progress for E799, the analysis presented in this thesis are fortunately not sensitive to these effects. On the other hand, delta ray production can give rise to low sod pairs since signals from a delta ray can arrive at the sense wire before the signal from the “real” track. This effect is simulated in the monte carlo.




Figure 12. The sum of distances, or *sod*, distribution for $K_L \rightarrow 2\pi^0$ events in DC2 in the E799 winter run. The mean of the *sod* distribution corresponds to $\frac{1}{2}$ a cell width in the DC, as seen in figure 10.

2.2.2.2 The Analysis Magnet

The magnetic field in the analysis magnet is oriented vertically so that the charged tracks bend in the x-view. The size of the magnetic field is about 2 kilogauss in the vertical direction over a gap of about 2 meters. The momentum kick in the x-view provided by the magnet is given by the field integral

$$\Delta p_{\text{mag}} = \frac{1}{c} \int B dz \quad (2.3)$$

In E799, Δp_{mag} was about 205 MeV/c and needed to be calibrated throughout the run using $K_L \rightarrow \pi^+ \pi^-$ decays. (In fact, the calibration of Δp_{mag} provides the basic energy calibration for the whole experiment.) Within a good approximation, the momentum kick can be considered as taking place at a single Z plane located at the center of the analysis magnet. The XYZ dependence of magnetic field was mapped with a Hall probe “ziptrack” system. The field map was used to calculate the field integral in Z and provide a coarsely binned (about 5cm square bins) map of the variation from average of the momentum kick in X and Y at the center of the magnet. The whole map is then scaled up or down based on the calibration mentioned above. Small residual fringe fields exist between DC1 and DC2 and between DC3 and DC4. These are corrected for assuming a constant kick

factor of about 1 MeV/c reasoned out from the non-planarity of two track events in the upstream and downstream pairs of chambers respectively.

The magnetic field was flipped⁴ periodically during the run to control for possible systematic effects related to magnet polarity.

2.2.2.3 Momentum Resolution

The momentum resolution of the KTeV spectrometer was measured to be [Hanagaki 98]

$$\frac{\sigma_p}{p} = 0.016\% \times p + 0.38\% \quad (2.4)$$

The constant term is due to multiple scattering, while the resolution gets worse at higher momentum. This is because the momentum kick, Δp_{mag} , becomes small with respect to the total momentum p of the track implying that the bend of the track in the X view becomes harder to measure. The term linear in momentum is thus a result of the finite position resolution of the detector.

2.2.2.4 DC1-OR Trigger Source

The X and Y plane pairs of DC1 and DC2 were used as trigger sources in Level 1. This was possible despite the long drift times associated with the drift chambers. Although the maximum drift time was about 200 ns, the track is always between two wires in adjacent planes so that at least one of the hits should arrive in less than 100 ns, so that the logical OR of the wires can be a relatively prompt signal [KTeV 207]. Furthermore, due to the non-linearity of the XT relations, the early hit always arrives within 80 ns. The advantage of this trigger source is that it provides a “telescope” effect with the V and V’ trigger hodoscope planes downstream of the DCs. The effect was to knock down the L1 trigger rate by a factor of about 2, even though the sources spanned 4-5 buckets.

⁴The spin rotator magnet polarity was also changed from time to time for hyperon analyses. This led to one of my favorite unlikely sentences, “Flip the spin rotator.” or some such permutation.

Detector Element	Z position (m) at Upstream Face	Relevant Aperture (X by Y)
DC1	159.42	1.30 m square
SA2	165.10	1.54 m \times 1.37 m
DC2	165.57	1.64 m \times 1.44 m
Analysis Magnet	170.01	2.90 m \times 2.00 m
SA3	174.00	1.69 m \times 1.60 m
DC3	174.59	1.74 m \times 1.64 m
SA4	180.00	1.75 m square
DC4	180.49	1.90 m square

Table 5. KTeV Spectrometer Z positions and Apertures. Only the inner apertures are given. The Z positions are at the upstream face of the detector.

To make the L1 sources, the wires in the X and Y views were logically divided into 16 “paddles.” One set of sources fired if there were ≥ 1 hits in the plane pair and the other fired if there were ≥ 2 hits in the plane pair.

2.2.2.5 The Spectrometer Anti System

A number of square aperture veto counters were located just upstream of each of drift chambers 2, 3, and 4. These spectrometer antis (SA’s) were made of Pb-scintillator sandwich and were, like the RCs, designed to detect stray photons down to energies of 100 MeV [KTeV 132],[KTeV 102]. Each SA had 16 radiation lengths of Pb like the RCs, but it was arranged in 32 layers of 0.5 radiation length Pb and scintillator. The modules of the SAs were beveled at an angle of 13.4° with respect to the beam to minimize cracks. These veto counters defined the apertures around drift chambers 2,3, and 4. Light was guided out of the SA modules by glass fibers to phototubes mounted around the outer edges of each counter. The phototube signals were digitized by LeCroy 4300 ADCs. Discriminated signals from the SAs also provided sources to the Level 1 trigger to quickly veto an event

Component	Thickness (or equivalent)	Radiation Length
Cathode Wires	1.26e-3 cm	0.06%
Anode Wires	4.05e-5 cm	0.01%
Mylar Windows	0.02 cm	0.04%
Radiator	5.2 cm	1.21%
Active Gas	2.5 cm	0.11%
Buffer Gas	1.0 cm	0.02%
Total		1.45%

Table 6. Breakdown of the radiation length of the one of the KTeV TRD chambers. The total is calculated for the in-acceptance part of the chamber. In the regions intersecting the neutral beams, there is no radiator material.

on the detection of decay products that would otherwise escape the detector. These sources were set to fire if any SA registered more than 0.5 GeV of energy and provided this information on a bucket by bucket basis.

2.2.3 The Transition Radiation Detectors (TRDS)

The transition radiation detectors (TRDs) provided efficient π/e rejection during E799. For E832, the TRDs were rolled out in an effort to minimize the in-acceptance material. The TRDs will be discussed at length in Chapter 4. The TRDs account for roughly 11.5% of a radiation length of material in the “in-acceptance” region (assuming 8 TRD modules.) The radiation lengths of the components of one TRD are shown in table 6. [KTeV 158]. The radiators had holes cut out in the beam region to reduce the possibility of beam interactions in the TRDs.

The TRD system also provided electron tagging capability at Level 2. The TRD trigger is described elsewhere and is not used in this analysis.

Detector	Z position (m)	Aperture
DC4	180.49 m	1.90 m square
TRD1	181.17 m	2.18 m square
TRD8	183.36 m	2.18 m square
V hodoscope	183.90	1.9 m square
V' hodoscope	183.95	1.9 m square

Table 7. Z positions and Apertures for the trigger hodoscopes and the TRD system. None of these apertures is limiting.

2.2.4 The Trigger Hodoscopes

A common requirement of all charged triggers in KTeV is the firing of the trigger hodoscopes. Located just downstream of the TRDs, these two planes of scintillator, labeled V and V' , provide fast pulses upon the traversal of charged tracks. Each scintillator plane was 1.0 cm thick and consisted of an array of paddles in the vertical direction. Each paddle spanned roughly half the distance from the top or bottom of the counter plane to the midline. This is illustrated in figure 13. Once again, holes are left in the beam region to minimize beam interactions. The paddles were not overlapping in each plane, so that cracks between paddles could become an efficiency problem. However, the paddles were carefully packed and tightened together in the plane by screws to minimize the cracks. Furthermore, the paddles were of different widths in each plane and were arranged so that there were no overlapping cracks between the two planes. The trigger sources are based on the discriminated signals from the paddles on V and V' . The signals are analog summed and discriminated again to provide the sources ≥ 1 or ≥ 2 or ≥ 3 hits in V or V' .

2.2.5 The Cesium Iodide Calorimeter

The CsI calorimeter provided precise measurements of the energies of incident electromagnetic particles, e^\pm and γ . Each incident e^\pm or γ deposited energy in an island of CsI




Figure 13. The V and V' trigger hodoscope planes. The paddles shown are of different widths so that there are no overlapping cracks when one plane is overlaid on the other.

crystals called a cluster. The depth of the CsI was optimized to contain completely the electro-magnetic showers up to 200 GeV. The energy resolution of the calorimeter is approximately [KTeV 403]

$$\frac{\sigma_E}{E} = 0.55\% + \frac{0.7\%}{\sqrt{E}}, \quad (2.5)$$

where the constant term comes from shower fluctuations and the second term comes from photo-statistics. This corresponds to about 20 photo-electrons per MeV. The excellent energy resolution of the CsI calorimeter has two important consequences for this analysis. Firstly, The CsI contains only 1.4 hadronic interaction lengths of material. Showers due to hadronic backgrounds, π^\pm , are thus not fully contained by the calorimeter. By measuring the ratio of energy deposit in the calorimeter to the momentum as measured by the charged spectrometer, one can effectively discriminate electrons from pions. The effectiveness of this form of identification improves with the energy/momentum resolution of the electrons. Secondly, the mass resolutions of various reconstructed particles improve with the energy resolution. Cutting on the deviation of the mass of the reconstructed 2γ system in the $K_L \rightarrow \pi^0 e^+ e^-$ analysis, for example, rejects the $K_L \rightarrow ee\gamma\gamma$ background. Since $K_L \rightarrow ee\gamma\gamma$ is essentially flat at the π^0 mass, the effectiveness of this cut is inversely proportional to the resolution.

Detector	Z position (m)	Aperture
V' hodoscope	183.95	1.9 m square
CIA	185.19	1.84 m square
CA	185.91	frames beam holes
CsI	186.01	1.90 m square
Pb Wall	188.53	(Not Limiting)

Table 8. Z positions and apertures for detector elements around the CsI. The limiting apertures are in **bold**.

2.2.5.1 The CsI Crystals

The CsI calorimeter was constructed from 3100 blocks of pure CsI [KTeV 403]. Each block was 50 cm long, corresponding to 27 radiation lengths or 1.4 hadronic interaction lengths of material along Z. There are 868 “large” crystals, each $5 \times 5 \times 50$ cm, and 2232 “small” crystals, each $2.5 \times 2.5 \times 50$ cm. Most of the native crystals were grown and cut to 25 cm lengths and then glued together to make 50 cm lengths. About 20% of the crystals were able to be grown at the full 50 cm length. The crystals were arrayed in a square array 1.9 meters on a side with 15 cm square holes to allow passage of the neutral beams. The large crystals populated the outer edges of the array while the smaller crystals populated the interior. The beam hole centers were at $Y = 0$ and $X = \pm 15$ cm. The CsI array is shown in figure 14.

Each crystal was wrapped in a custom arrangement of white reflective Teflon, reflective aluminized mylar, or black tape. The wrappings were tailored to better achieve uniformity of response along the length the particular crystal as measured with cosmic rays. Each crystal was attached to a photomultiplier tube with a digital base (DPMT). The optical coupling to the phototubes is provided by transparent RTV cookies. Also, a special filter was used in each crystal to allow passage of the fast (uV) component of the light to the phototubes but absorbed the slow component. Lastly, opaque filters were used on some crystals in order to tune the amount of light output so that the it was more or less uniform over the array.




Figure 14. The CsI calorimeter was constructed from 3100 crystals of pure CsI in a square array with two beamholes in the center to allow the neutral beam to pass through. The size of the array is 1.9 m \times 1.9 m. Two 15 cm square beamholes exist at $(X,Y) = (\pm 15.0\text{cm}, 0.0\text{cm})$ to allow for passage of the neutral beam. A 1.2 m decorated hardboiled egg is shown for scale.

2.2.5.2 *CsI Frontend Electronics*

The CsI phototubes were run at -1200 V and provided a gain of 5000. Peak current output was about 30 mA corresponding to about 80 GeV of energy. The DPMTs included elements to provide high voltage to the dynode stages of the phototubes, digitize the signal, and buffer the result. The QIE is the part of the DPMT circuit that takes an input current and presents a voltage to an 8-bit FADC for digitization along with an exponent that indicates the range of the input current. Current from each phototube was split into eight binary ranges, ie. $I/2$, $I/4$, ..., $I/256$. Then, on a per bucket basis, the charge in each range is integrated on a series of capacitors. Voltage comparators indicate which of the binary ranges is “in range,” and provide a 3-bit exponent. The in-range capacitor charge is then processed by an 8-bit FADC, the Analog Devices AD9002, to

provide a mantissa. The 8-bit mantissa with 3-bit exponent provides 16 bits of dynamic range, from a few MeV to over 60 GeV. The QIE was a custom built circuit developed at Fermilab. [KTeV 203]. The QIE operates with a 53 MHz clock. The charge digitization is a 4 stage process carried out in round-robin fashion on 4 identical circuits, each with its own set of capacitors, in each DPMT base, allowing for deadtimeless readout. To synchronize the results with the correct capacitors, a capacitor ID bit pair is included in the QIE output.

2.2.5.3 *The Hardware Cluster Counter (HCC) and Ettotal (Et) System*

In order to quickly count clusters in the CsI, KTeV instrumented a hardware cluster counter (HCC) [KTeV 298],[Bown 96]. The hardware cluster counter received discriminated and latched dynode signals from each of the 3100 CsI crystals, providing a map of the “on” channels in the CsI array. This information came every bucket, and was stored in a FIFO 2048 events deep. On receipt of a Level 1 trigger, the corresponding data is read from the FIFO into the HCC system. This information was analyzed in terms of cluster topology. The array was broken up into 2×2 crystal groupings, and each grouping was given a weight according to which crystals are “on.” The weights are listed in figure 15. “Left” turns were weighted by a -1, “right” turns were weighted with a +1, etc. The HCC summed the weights in the entire array. The sum of weights, when defined in this way, is equal to four times the cluster count.⁵ The output of the HCC was a four bit code. The lower three bits was a cluster count and the high bit was an overflow marker, so that the final result was a number from 0-7 or >8 clusters. The primary advantage of the algorithm is that it took a finite amount of time to sum the clusters in the array. The time to decision is 2 μ sec, so that the HCC was used in the Level 2 trigger. This processor was, nonetheless, the longest processor to return a decision. [Haney 92].

KTeV also instrumented a fast Ettotal (Et) System in order to obtain a quick estimate

⁵This is easily verified on a few examples of clusters. Simply put, it is the number of right turns minus the number of left turns. One example where the algorithm fails is the case of a “donut” cluster where many “on” crystals completely surround an isolated “off” crystal. The weight in this case is zero. However, this case never happens in normal operation.



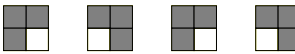


Hit Block Patterns	Pattern Value
	0 - No Turns
	+1 - One Right Turn
	-1 - One Left Turn
	+2 - Two Right Turns
	0 - No Turns

Figure 15. A diagram showing possible pattern inputs into the HCC and their weights. In the HCC, all 2×2 patterns of crystals are processed according to which crystals are above threshold. The sum of the weights taken over the entire array is 4 times the number of clusters.

of the total energy deposit in the CsI on a bucket by bucket basis. The Et system summed the energies in the 3100 CsI crystals also using the dynode signals. These were analog summed in stages, culminating in sums for the four CsI quadrants and finally for the whole array. This sum was then discriminated and provided sources to the L1 trigger system. There were four different Et thresholds provided, and the analysis presented here used a 25 GeV threshold on E_t . [KTeV 196].

2.2.5.4 The Collar Anti (CA) and the Cesium Iodide Anti (CIA)

At the CsI, there are two detector boundaries to worry about. The outer boundary of the CsI is guarded by the Cesium Iodide Anti (CIA). The CIA is alike in design to the SAs, and also provides a similar veto source to the L1 trigger system. The inner boundary of the CsI consists of the beam holes and is guarded by the Collar Anti (CA). The CA consists of 3 layers of tungsten-scintillator sandwich [KTeV 177]. Each layer of tungsten is 2.9 radiation lengths in Z. The CA overlaps the beam hole blocks of the CsI by 1.5 cm, and provides a veto source to the L1 trigger system. The CA is depicted in figure 16.




Figure 16. In this closeup view of the CsI beamholes, one can get a glimpse of the wily Collar Anti (CA). This beast lives in front of the beam holes of the CsI, overlapping the innermost layer of crystals by 1.5 cm. The CA guards the hole edges in the CsI and vetoes events which are too close.

2.2.6 The Lead Wall and Hadron Anti

Behind the CsI calorimeter, a 15 cm thick Pb wall was installed followed by a scintillator plane called the Hadron Anti (HA) [KTeV 191]. Hadronic showers initiated in the Pb wall were detected by the HA so that we could bias against events with π^\pm at the trigger level. The thickness of the Pb wall was chosen to absorb leakage of electromagnetic showers from the back of the CsI to be sure that signals in the HA were really initiated by a hadronic shower. Thus rejection of K_{e3} background was optimized while maintaining a high efficiency for the electro-magnetic modes. An analog sum of all the counters was discriminated and used in the level 1 trigger. The Pb wall and HA had one hole in the middle to allow the neutral beams to pass through. The size of the hole was 64 cm in X and 34 cm in Y. The signals from the 18 HA paddles were analog summed and the sum was discriminated. For the analysis in this thesis, the threshold was about 2.5 MIPS.⁶ This discriminated signal supplied a veto source to the L1 trigger. The HA is shown in figure 17, and the positions of various detectors behind the CsI are shown in table 9.

⁶MIP stands for *minimum ionizing particle* and refers to the amount of energy left in a detector by an energetic ionizing particle above the minimum of the Bethe-Block formula.




Figure 17. A drawing of the HA counter bank. The HA sat immediately downstream of the Pb wall and was used as a veto on the detection of hadronic showers at Level 1.

Detector System	Z Position
Pb Wall	188.53
HA	188.97
Filter Steel 1a	189.09 to 190.09
BA	191.09
Filter Steel 1b	191.74 to 194.74
MU2	194.83
Filter Steel 2	195.29 to 196.29
MU3	196.36
Hans Kobrak's Cabinets	198.0 to 199.0

Table 9. Detector elements downstream of the CsI with Z positions. There were no limiting apertures here. Only the HA is relevant to this analysis through the trigger, and MU2 likewise for part of the run.

2.2.7 And the Rest ...

2.2.7.1 90° Accidental Counter

An important system that hasn't been mentioned so far is the 90° accidental counter. The 90° counter was a small charged particle telescope that surveyed the target through a small hole in the side of the target pile. As its name indicates, it viewed the target along an axis perpendicular to the beam axis. The telescope was comprised of three scintillator planes. A coincidence of all three planes formed an accidental trigger. Since the telescope viewed a substantially different part of the target space, activity in the 90° counter was expected to be proportional to the beam intensity but uncorrelated with activity in the detector. Accidental triggers then provide a snapshot of accidental activity in the detector. These accidental events later formed the basis for the simulation of accidentals in the KTeV detector.

2.2.7.2 The Muon System and BA

The HA was followed by 4 meters of steel designed to range out remnants of hadronic showers, but pass muons of momentum $> 7\text{GeV}$. Behind this initial amount of steel was a scintillator hodoscope called the MU2 bank. MU2 consisted of a single plane of overlapping scintillator paddles oriented vertically. This bank was followed by another meter of steel and another scintillator hodoscope MU3 which consisted of two planes of scintillator paddles, one oriented along the vertical and the other along the X direction. The MU2 bank was hermetic and used to veto events with muons, and provided a source to L1 indicating ≥ 1 or ≥ 2 hits. MU2 was used for part of the winter run in the trigger, but this requirement was eventually dropped because the accidental rate was too high. The MU3 planes are used to trigger on events with a muon in the final state, and are not used in this analysis. In a gap in the first steel filter was the back anti (BA). The BA was designed as an active beam dump. It was not in the trigger and was not used in these analysis. The BA is described in [KTeV 187].

CHAPTER 3

Locating the Interesting Events

During the E799 run, the rate of K_L produced at the target was about 1 per bucket, or 53 MHz during the spill. However, the rate of K_L particles that actually decayed in our fiducial volume was only about 1 MHz. Also, neutron interactions led to a rate on the order of a MHz during the spill. Although an impressive data acquisition system (DAQ) existed, it was impossible to write out all of the events. In fact, the DAQ was capable of about writing 60,000 events to tape per spill. Two levels of hardware trigger and one level of software trigger were employed to knock down the event rates to more manageable levels. In addition to the DAQ, there was a slow-control/data acquisition system (SDAQ) for the monitoring of environmental data, voltages, and for detector control. After data collection, the set of events was further refined by a data split, which divided the raw data into subsets based on physics interest; and a data crunch, where the data in each subset were analyzed offline according to general criteria. These data refining techniques are the subject of this chapter. I will focus primarily on the requirements of the analysis presented in this thesis.

3.1 The Trigger

E799 was able to write a maximum of about 60,000 events to tape during a spill. In order to make sure that the tapes weren't being filled with uninteresting events, the trigger system used detector information that was quickly available to make an educated

guess about the nature of a given event and made a decision to either keep or reject an event. The first two levels of this trigger system were implemented in hardware and the third in software. At each stage, it was important to keep processing time and event rate at a minimum to minimize deadtime.

Level 1 (L1) used the fastest signals available and made a decision in each RF bucket. Therefore, L1 incurred no deadtime. The level 2 (L2) trigger also used quick signals from the detector, but involved significant processing time and produced more sophisticated information. The trigger system was dead during L2 decision making. The level 3 (L3) trigger was implemented in software and produced a full reconstruction of the events that passed level 2. During the on-spill period, events accepted by L2 were buffered and analyzed by the L3 system during both the on-spill and off-spill periods. If the processing finished before the arriving of the next spill, there was no deadtime incurred by L3. Overall, E799 was typically incurred a deadtime of about 33% with 4×10^{12} protons on target.

L3 had access to relatively imprecise “online” calibration constants, and therefore was capable of making only the most general purpose and loose cuts on the data. The data were further tagged at L3 as being important to one or another analysis.

Other trigger duties included the collection of calibration and pedestal data for the various detectors and the collection of “minimum bias” data, events with the loosest practical trigger requirements, for calibration and systematic studies. Data passing the L3 trigger were written out to tape. The name of the trigger used in this analysis is “2E-NCLUS” trigger, because we are interested in events with 2 electron tracks and $N \geq 4$ clusters.

More information about the trigger can be found in [Zimmerman 99], [KTeV 198], [KTeV 199], [KTeV 201], and [KTeV 361]. More information about the data acquisition system can be found in [KTeV 184], [KTeV 389], and [KTeV 369].

3.1.1 The L1 Trigger

The L1 trigger consisted of logical combinations of the fastest signals available. With the exception of the DC-OR signals, these came from phototubes : signals from the RC system, the SA/CIA system, the trigger hodoscopes V and V' , E_t from the CsI calorimeter, and the HA hodoscope. (These trigger sources are described in Chapter 2.) The phototube sources provided information on a bucket basis, while the DC-OR sources were spread out over 4-5 buckets.

The L1 trigger sources were delayed so that they arrived at the trigger system at about the same time. Sources from upstream detector elements like RC6 were delayed more than those from downstream elements to account for the time for decay products to physically traverse the detector. Even so, because of the finite size and signal propagation in some of the larger detector elements, this synchronization could only be done to within about 10 ns. Therefore, the L1 sources were resynchronized by requiring them to be coincident with an accelerator supplied 53 MHz bucket signal. After resynchronization, the sources were spread in time by only 1-2 ns.

The L1 sources were input into LeCroy 4413 memory lookup units (MLU). Since there were more sources than inputs on a single MLU module, the sources were divided into groups, each of which was serviced by its own MLU. Coincidence logic performed on the outputs of the MLUs by a pyramid of LeCroy 4508 programmable logic units (PLU) generated an actual trigger.

The trigger sources used in this analysis are

- **!PHV** : No RC, SA, or CIA module fired above 0.5 MeV threshold.
- **!CA** : No CA counter fired above 14 GeV.
- **!HA** : HA sum did not fire above 2.5 MIPS.
- **!MU2** : No hit in the MU2 hodoscope plane.
- **2V** : At least two hits in V and one in V' or vice versa. This essentially allows one hit to be missing because of cracks between V, V' paddles.

- **ET_THR3** : There must be more than 25 GeV deposited in the CsI as measured by the E_t system.
- **DC12** : At least one logical “paddle” hit in DC1 or DC2.

The L1 part of the 2E-NCLUS trigger is then

$$\mathbf{2E-NCLUS(L1)} = \mathbf{!PHV \cdot !MU2 \cdot !HA \cdot !CA \cdot ET_THR3 \cdot DC12 \cdot 2V}. \quad (3.1)$$

During a typical winter spill at 4×10^{12} protons on target, there were 1×10^6 L1 triggers of which 3.8×10^5 were 2E-NCLUS triggers. In the summer, the overall trigger rate was about the same, but there were about 4.3×10^5 2E-NCLUS triggers due to changes in the trigger definitions in the summer.

Small changes to the trigger occurred during the run for technical reasons. The CA was used only after run 8281 in the winter. Due to high accidental rates in the MU2 counter, the MU2 veto requirement was dropped after run 8577. About 50% of the data in the winter run had the MU2 veto in the trigger. Most of the run did not actually include the RC8 veto in the set of photon vetoes because of non-standard behavior. It was included briefly during the winter run with a high threshold.

3.1.2 The L2 Trigger

The L2 trigger system included six systems; hit counting in the DCs, Y-track finding in the DCs, in-time pair finding in the DCs, cluster counting in the CsI, electron tagging in the TRDs, and a stiff track trigger for hyperon studies. The decision times of these systems ranged from 800 nsec to 2.5 usec. The TRD trigger, stiff track trigger, and Y-track finder were not used in this analysis and are described elsewhere, [KTeV 213], [KTeV 213], [KTeV 361].

During L2 processing, the L1 trigger was inhibited. L2 took about 2.5 us to make a positive decision. If L2 accepted an event, an additional 15 usec of deadtime was incurred. However, L2 could also abort an event early based on one of the faster subsystems. The trigger was reset within 0.5 usec of issuing a L2 abort.

3.1.2.1 *Hit Counting*

Hit counting in the DCs served to reject events that did not contain enough hits to reconstruct the required number of tracks. The hit counting logic was implemented by the Kumquat modules designed at the University of Colorado. The number of hits was counted in each of the eight DC views, (X and Y views in 4 DCs.) The Kumquats could not tell if the hits were out of time, however. The processing time of the Kumquats was 800 ns. The hit counting requirement for this analysis was

- **2HCY_LOOSE** : At least 2 hits in each Y-view were found, but at most 1 missing hit in either DC1Y or DC2Y was allowed.

3.1.2.2 *In-time Hit Pairs*

Due to the long gate times needed in the DC system, reported hit pairs were not always in-time with the bucket that triggered L1. In order to reject out-of-time hit pairs, the hit pair times were measured with a TDC. The out-of-time hit pairs were rejected by the Banana modules, also constructed at the University of Colorado. To get an idea of what the Bananas do, we can plot the time of one of the DC hits, t_1 , against the time of the other hit, t_2 . The shape of the distribution of good in-time hit pairs in the t_1, t_2 plane resembles a banana; hence the name of the module. Out-of-time hit pairs also make bananas, but their bananas do not correspond to the correct sod mean. The Banana modules use a memory lookup to select the in-time banana. This is illustrated in figure 18. The Banana modules required 800 ns to make a decision. The Banana requirement for this analysis was

- **1HC2X** : At least 1 Banana accepted hit pair in the DC2 X-view was found.

The Bananas were installed and working after run 8262.




Figure 18. The selection of in-time hit pairs by the Banana modules. The figure on the left shows the selected in-time hit pairs which are shaped like a banana. The upper triangle above the banana consists of low-sod pairs which were often reconstructible offline; therefore these were also accepted at trigger level. The figure on the right shows the out of time bananas. The bananas could either be early (unripe) or late (spoiled.) Also, events with unphysical sods may not contribute to any of the bananas.

3.1.2.3 Hardware Cluster Counter

The HCC was described in chapter 2. It required about 2.2 usec to make a decision and was the last L2 processor to return a decision. The HCC requirement for this analysis was

- **HCC_GE4** : At least 4 HCC clusters were found.

3.1.2.4 2E-NCLUS L2 Trigger Requirements

The L2 part of the 2E-NCLUS trigger is then

$$\mathbf{2E-NCLUS(L2)} = \mathbf{2HCY_LOOSE} \cdot \mathbf{1HC2X} \cdot \mathbf{HCC_GE4}. \quad (3.2)$$

The total L2 trigger rate was about 2.0×10^5 triggers per spill. Of these, about 52000 were 2E-NCLUS triggers in the winter and about 59000 were 2E-NCLUS triggers in the summer.

3.1.3 The L3 Filter

Events that passed the L2 trigger were read out completely into four SGI challenge machines. The L3 filter was implemented in software and consisted of a simplified offline reconstruction of the events. Events passing L3 were written to tape. At this point in the data selection, events passing 2E-NCLUS at L2 were further classified according to more specific physics based criteria, and L3 assigned a tag based on this classification. In this analysis, the L3 tag of interest is the “2E-NCLUS” tag, and consisted of

- At least 2 tracks must be able to be reconstructed using loose tracking requirements.
- At least one good decay vertex candidate must be found.
- There must be ≥ 4 clusters in the CsI.
- Both tracks must point to clusters in the CsI Calorimeter.
- At least 2 tracks satisfy $E/p > 0.75$ to identify electron candidates.

An event had to satisfy requirements for at least one of the tags in order to be written out. In addition, a prescaled amount of random accepts were taken in order to study the efficiency of various elements of the L3 trigger.

3.1.4 Other Triggers

In addition to the 2E-NCLUS trigger described above, there were other special triggers of indirect use in this analysis. These included the calibration triggers :

- CsI calibration laser flash (5 Hz, on and off spill)
- CsI calibration cosmic ray¹ (30 Hz, off spill only)
- DPMT pedestals (3 Hz, on and off spill)

¹A cosmic ray telescope consisting of long scintillator paddles above and below the CsI made special triggers. Because the telescope was segmented in Z, measurements of crystal uniformity along the Z direction could be monitored throughout the run.

- Other pedestals (off spill only, every tenth spill)
- TRD Xray triggers² (10 Hz, off-spill)

In addition to these calibration triggers, a special “minimum-bias” trigger took physics data during the spill. This 2TRACK trigger consisted of

$$\mathbf{2TRACK} = \mathbf{2V} \cdot \mathbf{DC12} \cdot \mathbf{2HCY_LOOSE}. \quad (3.3)$$

The 2TRACK trigger was prescaled by a factor of 200 in the winter and 500 in the summer.

3.2 The Data Acquisition System

The Data Acquisition System (DAQ) moved ADC, TDC, scaler, and trigger information from the detector into the L3 processing memories and eventually onto tape. The actions of the DAQ system were orchestrated by the trigger system in the sense that no data was moved without at least a partial trigger decision.

3.2.1 The CsI Readout and Digital Pipeline

Signals from the CsI phototubes were digitized by custom QIE and FADC components on each of 3100 DPMT cards at the back of the CsI calorimeter, as described in chapter 2. This digital information was buffered by Driver-Buffer (DBC) chips on the DPMT cards. On receipt of a Level 1 trigger, 32 time slices of information was sent to a special holding buffer in the DBC. (A time slice corresponds to one bucket, and the 32 readout slices include information from before and after the in-time bucket.) On receipt of a Level 2 accept, the contents of the holding buffer were transferred to the digital pipeline.

The digital pipeline applies a threshold to the readout. The purpose of the threshold is to reduce the rather large amount of data that would be read out from 3100 crystals \times 32 slices. The threshold used in this analysis was 12 counts, corresponding to about

²These were of limited value because of the low energy of the calibration Xray.

7 MeV per crystal. Furthermore, the Pipeline could be configured to read out a subset of the 32 slices. For the analysis presented here, a total of 9 time slices of CsI were read out for crystals above threshold. The nine slices included three slices before the in-time bucket, the slice corresponding to the in-time bucket itself, and five slices after the in-time bucket. The pre-bucket slices could be used to reject early clusters, and the post-bucket slices were needed to reconstruct the cluster energies. In the following analysis, only four of the six slices will actually be used to reconstruct energies.

Upon readout request, the digital pipeline sent its data through CTIRC (Crate Trigger Interface/Readout Controller) modules to the L3 filter processors.

3.2.2 Other Readout

The readout of other ADC information, TRDs and RCs for example, occurred through the FERA (Fast Encode and Readout) ADC system. These were LeCroy 4300 ADCs which were read out through a custom Fermilab built DYC3 module. Like the Digital Pipeline for the CsI, the FERA ADCs could be used with a threshold to suppress readout of channels with no energy deposit. The threshold used here was the measured pedestal plus several counts. Other information, such as drift chamber TDC times and scalers, were also read out using the DYC3. The DYC3 was custom built by Fermilab for fixed target experiments and was part of the DART (Data Acquisition at Fixed Target) project. The DYC3s were CAMAC controllable modules that read data from CAMAC modules located in the same crate. The data was accessed through front panel interface (rather than the slower CAMAC bus) and assembled serially into “sub-events” using a token passing scheme within the crate. The sub-events were buffered in the DYC3 with a word count until readout.

Coupled with the DYC3 modules were CTI modules which had a bit mask from the trigger system which enabled the DYC3 to know whether or not the trigger was a calibration trigger or a physics trigger. It also contained an ordinal number code. The ordinal number was copied to all crates engaged in readout and was used to tag the the data being read out so that they could be assembled correctly into events later.

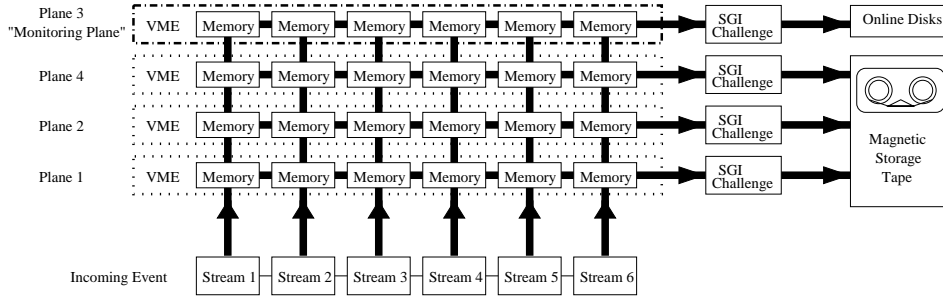


Figure 19. A depiction of the DAQ system. Event fragments traveled from the front end crates in 6 streams, each of which had 40 Mb/sec bandwidth. Fragments from a contemporaneous event were held in a single DAQ plane and analyzed by the L3 Filter. On a L3 accept, the event was to a contiguous block of memory and out to tape. The monitoring plane was the destination plane for monitoring events and calibration triggers. Calibration trigger data was written to tape from plane 3.

The DYC3 readout crates were arranged in two token rings. Another token passing scheme existed for readout within a token ring. As the token was passed around the ring, the DYC3 holding the token would pass the contents of its buffer to the L3 processors. If the DYC buffer became full, then it had the capability of inhibiting triggers until readout was complete.

3.2.3 Streams, Planes, and Event Building

The information from the DYC3 modules and the CTI modules entered the event builder on six streams of data. Two streams of data correspond to the two FERA/DYC3 token rings, and the other four streams of data come from the CsI digital pipeline. On receipt of a L2 trigger, the sparsified detector information is sent out on the streams. As data comes in from the streams, they are assembled into events on one of the three DAQ planes used for L3 filtering.

Each DAQ plane consisted of VME dual-ported memory with I/O controllers. There were 4.6 GBytes total memory in all of the DAQ planes. Each memory node represented the intersection of one stream and one DAQ plane, as depicted in figure 19. A component of each node was responsible for building a table of VME memory addresses versus event numbers so that L3 could find all the pieces of a single event. Event numbers were generated by the trigger system and were matched to events fragments using the CTI

ordinal numbers. The DAQ planes could buffer events from one entire spill. There were four DAQ planes total. Three of these planes were used for events to be processed by level 3 trigger, and the last plane was used to buffer online calibration and detector monitoring events. The monitoring events were prescaled copies of events from other planes, so that there was no event loss due to monitoring. The planes were numbered 1-4, and plane 3 was the monitoring plane. Data from each plane was processed by its very own SGI challenge machine which shared memory with the VME nodes in its plane. The SGI challenges processing planes 1,2, and 4 had eight 200 MHz processors running the L3 trigger. The SGI challenge corresponding to plane 3 had ten 150MHz processors. Superimposed on the one-minute spill rhythm provided by the accelerator (see chapter 2) was the action taking place in the VME memories. An interesting dance of VME memory buffers being filled during the on-spill period followed by the eventual finishing of L3 processing during the off-spill period was displayed in the control room as a pair of moving read-write pointers, oddly reminiscent of the Omar Khayyam's moving finger.

On a successful L3 tag or random accept, the event pieces were copied written to a contiguous portion of memory and written out to of four DLT-III tape drives attached to that particular SGI challenge. Events were written in essentially round-robin order. Although some attempts were made to route events with certain features to certain planes, this sorting effort turned out not to be particularly useful. (Events were sorted offline later during the "split".) Each DLT-III tape could hold 15 GB of raw data with compression. We could run for about 8 hours this way before the tapes filled up.

3.2.4 The Slow DAQ : Monitoring and Control

KTeV monitored everything in sight. In addition to the data acquired directly from the KTeV detectors through the front end system, certain detector environmental and control data were reported to the control room on a minute basis by the slow DAQ. The slow DAQ reported temperatures and humidity in the CsI blockhouse, voltages in the drift chamber spectrometer and the TRD system, and the nominal operating voltages of every NIM and CAMAC crate in the detector hall. The slow DAQ was also responsible

for command and control functions in the KTeV detector, and for logging performance data to a slow DAQ database, based on the hepdb system of CERNLIB. Among other things the slow DAQ was responsible for controlling and logging voltages in the TRD system as described in the next chapter, and for interfacing with the EPICURE system that provided information about the beamline. The SDAQ is documented in [Alavi 99a].

3.3 Further Refinements of the Data Set

3.3.0.1 *The E799 Data Split*

E799 wrote a total of 850 DLT-III tapes in the summer and winter runs. Each tape had a mix of events that needed to be sorted into categories for analysis. Obviously, 850 is too many tapes to handle on a regular basis, so the events were sorted into subsets based on physics interest. There were 19 different output streams from the split. Some events had multiple physics interest, and therefore were duplicated in several different streams. Since the analysis presented here only used one of the output streams, this was not a concern.

The output stream of interest, the “2ENEUT” split stream, included all of the 2E-NCLUS triggers. The output of this stream contained 98 DLT-III tapes.

3.3.0.2 *The 2E-NCLUS Trigger Crunch*

The number of tapes at this stage is still cumbersome. Therefore, a crunch was performed on the 2ENEUT tape set. The crunch performed a slightly more sophisticated analysis on the events and therefore needed access to more sophisticated calibration constants. When these constants were available, the crunch was able to begin.

The analysis presented here made use of a single crunch stream. The data for $K_L \rightarrow \pi^0 e^+ e^- \gamma$ and the $K_L \rightarrow 2\pi_D^0$ normalization mode were contained in a five cluster subset of the crunch. The requirements of the five cluster “EEGGG” subset are :

- The 2E-NCLUS tag from L3 must be set. (This is to prevent volunteer events that failed the L3 trigger but somehow passed the other crunch requirements.)
- The tracking code must be able to find a vertex. The code is run in a loose mode in which a single Y-track is allowed to be shared among two X-tracks.
- The tracks must match to clusters in the CsI.
- E/p of the tracks must be >0.9 .
- There must be exactly 5 clusters as found by the HCC.
- The vertex must be in one of the two beams. Specifically, the extrapolation of a vertex along the target-vertex vector to the CsI face must fall within one of the two beamholes. The cut is loose so that the extrapolation can hit the CsI by as much as 5 cm.
- The invariant mass of the 5 cluster “EEGGG” event must be greater than 0.380 GeV.

After the data crunch, there are 3 DLT tapes of “EEGGG” output.

The exact definition of the data set can vary from analysis to analysis. The intersection of functioning detector systems and analysis requirements changes both with respect to time and with respect to physics interest. The evaluation of functioning detector systems takes place in KTeV on a spill by spill basis. Each spill was assigned a 32-bit mask in which each bit stands for a given problem. For the purposes of this thesis, a spill was deemed a “good spill” unless one or more of the following spill bits was flagged.

1: Trigger Problem

2: DPMT Problem, DPMT pedestal exponent is > 0

3: DPMT Problem, Bad DPMT Capacitor

4: DPMT Problem, Blown Comparator

5: DPMT Problem, dead DPMT (yeah, really!)

- 6:** DPMT Problem, DPMT ped drift
- 7:** DPMT Problem, DPMT gain drift
- 8:** DPMT Problem, Broken Dynode
- 9:** Digital Pipeline Problem
- 10:** CsI Problem
- 11:** Etototal System Problem
- 12:** FERA ADC Problem
- 13:** Drift Chamber Problem
- 14:** Photon Veto Problem
- 15:** Trigger Hodoscope Problem
- 16:** Muon System Problem
- 17:** HCC System Problem
- 18:** Kumquat/Banana System Problem
- 21:** DAQ/L3 Filter Problem
- 22:** Non-E799 run or Specialty Run
- 23:** Short Run That Ended Badly
- 29:** Beamline System Problem

The winter run suffered from a swapped HCC cable through the beginning portion of the run. The problem was fixed after run 8245 of the winter. This problem is flagged as bit 17 which is included in the list of reasons to reject a spill. Since the problem is modeled in the Monte Carlo simulation, we can still use this data; however, we don't check bit 17 before run 8245 to avoid losing this data. It was later verified that there were no other HCC problems to worry about for this data.

The MU2 veto requirement was dropped after run 8577, so bit 16 was not checked after run 8577.

To determine the efficiency loss due to bad spills, we study the losses in fully reconstructed $K_L \rightarrow 2\pi_D^0$ decays. Table 10 recounts the percentages of different bad spill reasons occurring in the periods of running after run 8245. The overall efficiency loss (post-8245) due to cutting on bad spills is 10% in the winter and 5% in the summer. Note the high number of events in the summer flagged by problem number 3. This is because of an error in the database. While we are waiting for this problem to be fixed, we do not cut on problem 3 in the summer.

3.4 Conclusion

In this chapter, we have gone from 850 DLT tapes of raw output to 3 DLT tapes of crunch output. This allows us to go hog wild on analysis, since the entire data set can at this point be analyzed in a few hours.

Code	Winter	Summer
TOT	27225	18167
OK	24539	17273
1	0	0
2	236	226
3	110	3762
5	138	33
9	2	1
10	0	0
11	0	0
12	13	1
13	0	0
14	0	0
15	0	0
17	0	94
18	0	0
21	0	0
22	0	488
23	0	0
29	0	0

Table 10. Bad Spills for the Winter and Summer E799 runs. The numbers in the table represent fully the number of fully reconstructed $K_L \rightarrow 2\pi_D^0$ decays with the bad spill reasons shown at the left. The data in the table are compiled post-run 8245 only. Before run 8245, every event was flagged by problem 17 due to HCC cable swap. There were 2230 reconstructible $K_L \rightarrow 2\pi_D^0$ events with this problem representing about 7.5% of the winter data. Furthermore, we do not cut on problem 3 in the summer because for the majority of cases, the presence of problem 3 reflects an error in the database rather than a true problem with the DPMTs.

98/09/15 18.00




Figure 20. Bad spill reasons in E799 from reconstructed $K_L \rightarrow 2\pi_D^0$ events. Bad spill reasons 24,26, and 28 correspond to TRD problems and are not checked in this analysis.

CHAPTER 4

The Transition Radiation Detector in Detail

The KTeV transition radiation detector (TRD) consists of eight multiple wire proportional chambers designed to measure X ray production by electrons in polypropylene radiators mounted in front of each chamber. In E799 phase I,¹ the immediate predecessor to this experiment, the search for $K_L \rightarrow \pi^0 e^+ e^-$ was hampered by backgrounds with final states including single charged pions. The semileptonic kaon decay, $K \rightarrow \pi^\pm e^\mp \nu$, when accompanied by some combination of radiated γ and/or accidental γ , can fake the $K_L \rightarrow \pi^0 e^+ e^-$ final state if the charged pion is misidentified as an electron. This background from this mode in E799-I required draconian kinematical cuts, resulting in a loss of efficiency for $K_L \rightarrow \pi^0 e^+ e^-$. In particular, a final kinematical cut made against $K \rightarrow \pi^\pm e^\mp \nu$ in [Harris 94] sacrificed 15% of the remaining $K_L \rightarrow \pi^0 e^+ e^-$ acceptance in order to root out only 35% of the remaining $K \rightarrow \pi^\pm e^\mp \nu$ background; and also, 85% of the total remaining background of 2.1 events was accounted for by $K \rightarrow \pi^\pm e^\mp \nu$.² In E799-II, the philosophy has been to avoid making such cuts by augmenting the electron tagging capability of the KTeV detector. While improved pion rejection was expected from the CsI calorimeter, the design goal of the KTeV TRD system was to supplement this with

¹E799-I was the immediate predecessor to this experiment, and ran in 1993. E799-II is a part of KTeV.

²Sometimes, the decay mode $K \rightarrow \pi^\pm e^\mp \nu$ is denoted as Ke3.

a 50:1 rejection factor of charged pions at 90% electron efficiency. Using the fact that E799-II has amassed a factor of 10 in flux over E799-I, and assuming an improvement in pion rejection from the CsI of 3-4, then a rejection of about 50 is indeed required to reduce the $K \rightarrow \pi^\pm e^\mp \nu$ background to a level of 0.1 events or better in E799-II.

In fact, in the standard KTeV $K_L \rightarrow \pi^0 e^+ e^-$ analysis, the number of $K \rightarrow \pi^\pm e^\mp \nu$ backgrounds before a TRD cut is estimated to be 2.0. This extra improvement is attributed to a cleaner beam environment than that in E799-I, and thus freer of accidental photons also needed for $K \rightarrow \pi^\pm e^\mp \nu$ to fake a $K_L \rightarrow \pi^0 e^+ e^-$ event. Thus in the $K_L \rightarrow \pi^0 e^+ e^-$ analysis, the TRD will be used to generate a cut that is 96% efficient to accept single track electrons and rejects single track pions at the level of 25:1. With this cut, the dominant background problem becomes $K_L \rightarrow e^+ e^- \gamma \gamma$. The TRD was not used explicitly in the $K_L \rightarrow \pi^0 e^+ e^- \gamma$ analysis, however.

4.1 Overview of the KTeV TRD System

The KTeV TRD system consists of eight TRD modules located between $Z=181.2$ and $Z=183.5$ m from target, as depicted in figure 21. Each module consists of a radiator, which produces X-rays upon the traversal of an energetic electron, and a Xe-filled, dual planed, multiwire proportional counter (MWPC) to detect these X-rays, for 16 total planes. The TRDs are mounted on rollers so that they could be rolled in and out of acceptance depending on whether or not KTeV was running in E799 mode or E832 mode. The roller system is depicted in figure 22. When the TRDs are in acceptance, they do not form a limiting aperture. The TRD system was designed with the constraint that the in acceptance radiator material should be less than 10% of a radiation length of material.

4.2 Optimization of the TRD System with Simulation

Transition radiation is produced when a fast, charged particle traverses the boundary between two media with different dielectric properties. The probability of emission per boundary is small, so it is advantageous to use a radiator consisting of many boundaries.




Figure 21. The TRD system is depicted along the beam axis. All 8 TRDs lie between the last drift chamber, DC4, and the trigger hodoscope planes. A spare TRD module, not shown, was also in place during the summer run immediately downstream of TRD8.




Figure 22. The TRDs are shown in cross section to the beam. During E832, the TRDs could be rolled out of acceptance. Flexible cable trays were used so that no special cabling operations were necessary during rollout or roll-in.




Figure 23. The conceptual design of a KTeV TRD module. Xrays generated in the radiator must pass through a double window buffer volume containing an 80/20 mixture of C₂F₆/CO₂. The width of this region is up to 1 cm. The Xrays are then detected in an active region containing an 80/20 mixture of Xe/CO₂. The active region contains the MWPC, consisting of two anode planes and three alternating cathode planes. The anode - cathode gap is 6mm, and the cathode wire pitch is 5 mm. The anode wire pitch is twice the cathode wire pitch, and the two anode planes are staggered from each other.

The radiation of interest is in the X-ray region, involving photons in the range of 2 to 50 keV. A simulation based on calculations in [Durand 75], [Artru 75], and [Cherry 74] was used to check the design parameters of the KTeV TRD system. To model the X-ray absorption in the chambers, fits to empirical X-ray absorption data appearing in [Storm 70] were used.

The main characteristics of transition radiation include a strong dependence of the total yield on the Lorentz factor, γ , below the saturation value, γ_s . For the KTeV radiators, the value of γ_s was about 8000. Above γ_s , the total yield is independent of γ . The average momentum of individual charged tracks in KTeV is about 20 GeV. While this varies from decay mode to decay mode for electrons and pions, it is clear that the electrons will be above the saturation value while the pions (and muons) are well below. Thus, the TRDs act like a threshold detector in γ .

The main characteristics of the MWPCs were fixed by mechanical considerations and by minimization of material in between the radiator and the TRD active volume. The total number of TRDs was optimized by simulation. Operating under the constraint of 10% of a radiation length of radiator material for the whole system, the total number of TRD modules was varied. Figure 24 shows the result of this optimization. The estimated rejection factor of the TRD system increases with the number of chambers, but then plateaus at about $N=8$. This can be explained in terms of the X-ray self absorption of the radiators. At $N=8$, the thickness of the radiator becomes comparable to a single X-ray absorption length at 10 keV, the peak of X-ray output predicted in the model. The thickness of each KTeV TRD radiator is 5.25 in, while the X-ray absorption length at 10 keV is 5.7 in. For N less than this, the radiators are thicker; and transition radiation generated at the front portion of the radiator becomes increasingly absorbed in the rear.

The energy deposited in the TRD then includes X-rays from electron transition radiation and conventional dE/dx ionization energy loss from any charged particle. The energy spectrum for pions (and muons) will be a single ionization peak while the energy spectrum for electrons will be the sum of the X-ray transition radiation summed with the ionization energy.

4.3 The Radiator

The classical radiator design calls for a stack of regularly spaced foils stretched in air. Instead of attempting to construct a foil stack of roughly 2.1 m square aperture, the radiator was made from stack of fiber blankets consisting of polypropylene fibers




Figure 24. For a constant amount of radiator material, the number of TRD modules in the system was varied and performance modeled by monte carlo. The estimated pion rejection plateaus around $N=8$. Below this value, the pion rejection falls rapidly. The radiator thickness at $N=8$ corresponds closely to a single X-ray absorption length at 10 keV. For smaller values of N , the radiators are thicker than this value, and X-rays are lost to radiator self-absorption.

of 17 micron diameter and 2.5 lb/ft^3 bulk density. Each radiator stack was 5.25 in thick. The overall thickness of the radiator was chosen corresponds to 1 X ray absorption length at 10 keV based on the above optimization. The fiber mat used in KTeV was based on one used successfully in balloon borne cosmic ray experiments conducted by D. Mueller. Nonetheless, many different kinds of radiator material were tested during a beam test with a prototype X ray detector (MWPC) at KEK Japan National Accelerator laboratory in a beam of 5 GeV/c electrons, including the sample provided by D. Mueller. Among the tested materials were several different kinds of polypropylene fiber mat and polypropylene foams. All of the different radiators were tested at constant cross section of 1 X ray absorption length at 10 KeV. The null case was a machined plate of polypropylene plastic of same material thickness. The sample provided by D. Mueller was chosen because it had the highest transition radiation yield [KTeV 146].

Once chosen, the most critical aspect of radiator quality was the purity of the fiber blankets. After an extensive search of suppliers, Buffalo Felt of Buffalo, NY was chosen to manufacture the KTeV radiator fiber blankets. Manufacturers of such fiber blankets often add whiteners to the finished product, such as TiO_2 . This would be the kiss of death to X-rays. Therefore it was explicitly required that the finished product be free of all additives and/or machine oil byproducts of the manufacturing process. As a final check, the X-ray absorption length of the delivered radiator material was tested on receipt. The purity of the delivered product measured in this way actually exceeded slightly that of the test sample.

The radiator blankets were affixed to a frame of 2.08 m square inner aperture. The frame was constructed of light plastic honeycomb clad with an eighth inch thick aluminum skin. An aluminized mylar window of 2 mil thickness was stretched across the frame before radiator stacking. Aluminized mylar was used because of it resists permeation of water vapor and other contaminants better than plain mylar (see table 11) into the radiator and MWPC volumes. Each radiator blanket was cut to size using a stencil and hot knife. In addition to the outer dimension of each blanket, two rough cut regions were cut away in the center of each blanket with an aligned stencil designed to allow the two KTeV neutral beams to pass through without interaction in the radiator. To monitor the final positions




Figure 25. The radiator frame is shown in cross section to the beam and along the beam direction. The frame itself is made from a light core of plastic honeycomb clad with eighth inch thick aluminum plate. The frame was stretched with a 2 mil aluminized mylar window to impede movement of water vapor across the boundary. The radiator stack was made from polypropylene felt blankets and were hung from pegs at the top of the frame. In addition, the radiator was glued down at the edges. Finally, holes were cut in the radiator in the beam regions to allow passage of the two KTeV neutral beams. The aluminum coating was etched away from the mylar around the beam holes so that the final position of the beam holes could be monitored after the radiators were mounted to the TRD chambers.

of these radiator beam holes after the radiators were hung in their final positions, the aluminized coating was chemically etched away around the edges of the beam holes.

In order to build each radiator to the desired thickness, 21 blankets were stacked on each frame. These were mounted on pegs at the top of each frame and further glued around all edges to provide mechanical support. After fabrication, the radiators were hung in upright position for about 1 week each to allow the radiators to sag into their final configuration. After this hanging period, the beam holes were enlarged to their final size using a hot wire and an aligned lucite template. The final size of the radiator beam holes allowed clearance of the KTeV neutral beams at the Z position of the TRD system.




Figure 26. A diagram showing the components of a TRD frame. There were three layers in each lamination. The middle layer was a single piece steel core depicted at left. The outer layers were each four pieces; G10 on top and bottom, and aluminum on the sides. The medium separating the lamination layers was provided by a thin sponge material and glue. When pressed together between two externally aligned flat surfaces and allowed to dry, a flat wire plane resulted. The front of the wire plane had G10 pieces clad with copper traces to provide mount points for wires to be stretched later.

4.4 The X-Ray Detector

The conceptual design of each TRD module is shown in figure 23. The MWPC has two anode (sense) planes and three alternating cathode planes. Between the radiator and the chamber volume lies a gas buffer volume.

The primary mechanical support of each TRD chamber was provided by an aluminum strongback of 6" wide aluminum square hollow tube. (See figure 28.) The inner aperture of the strongback was 2.13 m, well outside of acceptance. Each wire plane was made of a laminated steel and G10 fiberglass sandwich. The wire to wire gain of each channel in a MWPC is a sensitive function of the wire spacing in each drift cell. The anode to cathode spacing is provided by the thickness of the wire planes. In order to ensure the uniformity of gains of the wire planes to within 10%, the planes were constructed on a special tooling table that was flat to within a few mils. First, an outer square layer was arrayed on the table (see figure 26.) Glue was applied to the top of this layer with thin sponge strips to provide compressive resistance. Each outer layer consisted of two rectangular sections of G10 (top and bottom) and two rectangular sections of steel (left and right) laid out in a square. On top of this first layer of gluey sponge, a steel core was laid. Another gluey

Gas	Plain Mylar	Aluminized Mylar
H2O	5900	330
CO2	19	1.26
O2	2.9	-
N2	0.54	0.14
Xe	0.095	-
C2F6	< 0.0076	-

Table 11. Permeability of several gases in the KTeV TRD system across 2 mil mylar windows, aluminized or not aluminized. The units are in $\frac{cc}{day \times 100 in^2 \times atm}$. (Based upon a table given in [KTeV 154].)

sponge layer was applied on top of the steel core, and finally the second outer layer of two G10 rectangles and two steel rectangles was laid. Compressive pressure was applied to the final assembly by a flat top plate and series of clamps. The thickness was maintained by many precision machined spacers placed at various intervals around the inner and outer apertures of the final assembly. When the clamps were tightened, the plane thickness was thus fixed by the spacer thickness until the glue dried some 24 - 48 hours later. The G10 pieces on which wires were to be mounted had circuit patterns etched in copper into their outer surfaces for electrical connections. The thickness was checked after the gluing procedure at various points with a micrometer, and the wire plane was considered a reject if the variation in these measurements had a spread of more than 6 mils. The frames had holes at various points through which tightening screws were put later during the final construction phase.

In addition to the wire planes, the MWPC needs windows to contain the chamber gas. The gas windows were made of aluminized mylar and were stretched across thick aluminum square frames. Aluminized mylar was used to reduce permeation rates of impurities and chamber gas across the boundaries. Table 11 shows the permeation rates across mylar versus aluminized mylar windows for various gases. The upstream entrance to the active volume was special for two reasons, however. To achieve high X-ray detection efficiency between 5 and 25 keV, the main chamber active volume is filled with an 80/20

Material	L_x at 5 keV (m)	L_x at 10 keV (m)	Density (g/cc)
Air	0.215	1.707	0.0012
CO2	0.129	1.086	0.0019
C2F6	0.030	0.250	0.0057
Xe	0.003	0.011	0.0059
80/20 Xe/CO2	0.004	0.014	-
80/20 C2F6/CO2	0.036	0.295	-
Mylar	0.000260	0.002	1.39
Polypropylene	0.000680	0.006	0.9
Polypropylene Felt	0.027	0.145	0.04

Table 12. The X-ray absorption lengths, L_x , of various materials are presented with material densities. The X-ray absorption lengths are calculated from fits to data in [Storm 70].

Xe/CO2 gas mixture. Since the Xe/CO2 mixture is heavy and since the TRD aperture is so large, hydrostatic forces acting in the chamber volume tend to cause the window to bulge. Table 12 shows the densities of the gases of interest. The window bulging acts to change the efficiency of the MWPC as a function of position along the surface. Therefore, the upstream window was a double window made from 2 mil aluminized mylar of about 1 cm thickness.³ In this buffer volume was filled a heavy but X-ray transparent gas mixture of 80/20 C2F6/CO2. This gas mixture approximates the density of the active volume gas and thus balances the hydrostatic bulging of the inner window. The window positions are also sensitive to slight changes in pressure across their boundaries. In order to minimize this effect, the windows were stretched to a high tension of 20 N/cm. The gas system, to be described extensively in [Quinn 99], controlled the pressure differential to within ± 2 mTorr. Prototype measurements showed that this design was sufficient to control the position of the inner window to ± 50 microns [KTeV 185]. This is small compared to the X-ray absorption length of Xe even down to 5 keV.

³The thickness of the window plane is only 6mm at the edges; however, the outer window bulges into the air so that the thickness of the gap may be as large as 1 cm at some places.

The inner aluminized window surface was designed to hold voltage during normal operation. The purpose of this voltage is to activate the gas region of about 1.2mm between the upstream cathode plane and the window. If a voltage is not applied across this region, then charge deposited in this region does not drift to the anode and is not detected; so that as far as X-rays are concerned, the region would act as an opaque screen and lower the overall detection efficiency. It has been shown in the prototype that the normal excursions of the position of the inner window do not result in gain shifts larger than 1% [KTeV 185].

Upon completion of a TRD dual window buffer volume, it was filled to overpressure on some bottled gas. The pressure was monitored for several days to check for leaks. If no leak was found, it was considered a finished product.

The active region of each TRD is 1.76 m. This is the aperture over which the TRD wires were instrumented and read out. The wires were strung along the vertical to take advantage of the larger spread in charged track illumination in that view because of the analysis magnet. The wire pitch was 0.5 cm for the anode (sense) wires and 0.25 cm for the cathode (field shaping) wires. The wires used were 1 mil gold plated tungsten wires for the anode planes and 4 mil silver plated BeCu wires for the cathode planes. The wire frames were wound at lab 6 at Fermilab. The frames were rotated on a large paddle with the appropriate wire coming off of a spindle. A stepper motor moved the spindle on each rotation to provide the final pitch of the wire planes. Once wound, the wires were soldered to the copper traces on the G10 portions of the wire frame. In order to keep the final ADC channel count of the TRD system under 2000 channels, the sense wires in the central 64 cm of each sense plane were ganged together in pairs and the remaining regions were ganged in groups of four. The total channel count was 56 gang-4 channels and 64 gang-2 channels per anode plane. For all 16 planes, this leads to 1920 channels of ADC total.

In the regions of the TRD intersecting the neutral beam, the anode wires were deadened to reduce accidental activity on those wires. The deadening process was to thicken the wires in that region to reduce the gain there. This was achieved with electrodeposition on the wires in the beam regions. The electrodeposition fluid was 8% Copper Sulphate




Figure 27. The dual electroplating boat used to deaden wires in the KTeV TRD beam regions. The dead regions were large enough to contain the KTeV beams (E799 summer large beams) within a 3mm tolerance for all TRDs.

(Sifco Dalic 2050 Copper Acid) in solution. The solution was placed in a lucite double “boat” of shape and size necessary to contain the beam region for each TRD. The boat was overfilled just enough for the solution to form a meniscus. The sense plane was then lowered onto the boat so that the anode regions to be deadened lay within the meniscus of the solution. The wires to be deadened formed the negative electrode while the other electrode lay at the bottom of the boat. In order to bring the diameter of the wires from 1 mil to 3 mils, 7.5 mA per linear inch of wire in the bath was applied for 87.5 seconds and then 12.5 mA per linear inch of wire in the bath was applied for an additional 87.5 seconds. This process was used in a small test wire plane in which the thickness and quality of the plating was checked with a ruled magnifying eyepiece. Later, the test wires were used in a small prototype chamber with anodes operated at 2600 V in 80/20 Ar/CO₂. In this setup, the 5.9 Fe55 X-ray produced a scope trace of 150 mV in the live region and less than 4 mV in the dead region.

The final fabrication of the TRD MWPCs took place when the windows and wire frames were assembled to the aluminum strongback. The gas seal between the planes was common RTV sealer. When all of the frames and windows were in place, a final aluminum square tube backframe was applied and screwed down. Each screw was tightened into final position with a torque wrench to 20 ft-lbs of pressure. It was important that all screws had the same tightness to ensure that the final MWPC was not lopsided. The wire frames and windows were internally aligned to two pegs on the strongbacks. The positions of the pegs were known and related to external fixtures for final alignment. Before each chamber was closed up, a survey team optically checked the positions of selected wires

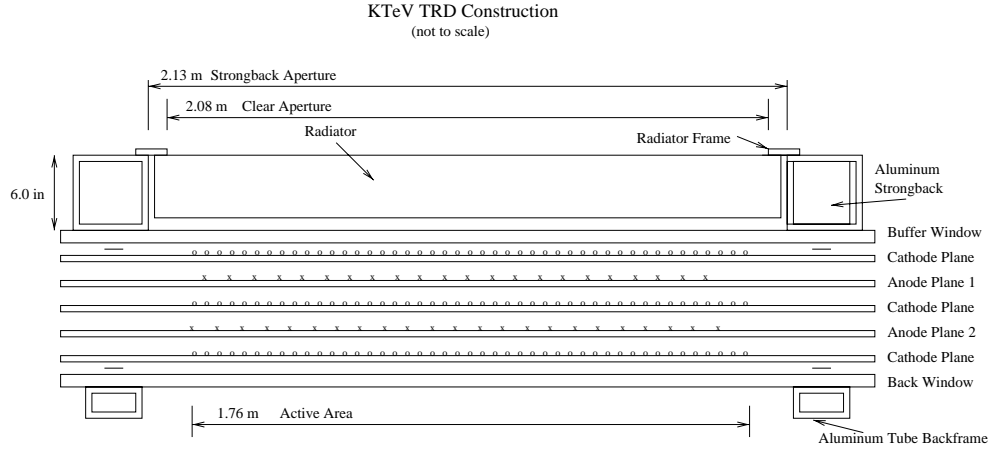


Figure 28. A diagram showing the major steps in the construction of the TRD chambers. The primary mechanical support comes from an aluminum strongback, which is 6" aluminum square hollow channel. An aluminum frame on which is stretched the two aluminized mylar windows comprises the C2F6 buffer volume. The steel-G10 laminated wire frames, both holding the wires and providing the anode-cathode separation distance, are stacked next. Finally, the rear window is added followed by an additional aluminum tube support. These elements are sealed with RTV sealer and screwed down to uniform tolerances.

in each plane to catch any gross misalignments. For the cathode planes, these wires were selected anywhere in the plane. For the anode planes, the thickened dead wire regions were sighted upon. Thus, the alignment procedure checked the final positions of the dead regions as well.

The finished chambers were allowed to dry for 24-48 hours and were then started on a bottle mixture of 80/20 Argon/CO₂. Once dry, testing of each chamber began.

Further details of the KTeV TRD MWPCs are found in [KTeV 185].

4.5 Testing and Commissioning

The strung frames were assembled with the aluminum strongback in a clean room at Lab 7. After each TRD completed assembly, it underwent a battery of tests with an Fe55 5.9 keV X-ray source. Typically, 80/20 Ar/CO₂ bottled gas mixture was started on the TRDs as soon as the gas fixtures were in place. After gas had been flowing for at least 24 hours, but after the RTV was judged dry enough, tests could begin on the chambers. The RTV was judged dry when the solvent could no longer be detected nasally at the

gas output. If the TRD passed all of the tests, it was warehoused in New Muon until the space was ready in KTeV hall. If the tests uncovered a problem, a determination was made to either live with the problem or to open the chamber and try to fix the problem. The tests included

- High Voltage Test. Using the CAEN high voltage supply, each chamber was slowly brought up to full voltage. If the chamber could not hold voltage and if the problem could not be tracked down to an external problem with the HV, then the chamber was reopened.
- Channel and Gain Test. Using the CAEN high voltage supply, the chamber was brought up to full voltage. Each channel were checked with a source to determine if it was live or if it had some other problem. If a channel was dead and it was in the central gang-2 region, then the chamber was opened.
- Dead Region Test. The collimated Fe55 source was shined on the deadened wire region of each plane. Using an oscilloscope and a tape measure, the dead regions were mapped to make sure that the correct regions were deadened.
- Gas Test. After the above tests, each chamber was run for an extended period of time on the gas recirculating system. When a new chamber was added to the system, the rates of gas impurities and the rates of gas loss from the system were monitored using the gas chromatograph. If a leak was found that was not external, then the chamber was opened and repaired.

During construction, only chambers 4 and 6 were re-opened because of problems found during testing.

After testing was completed and before final commissioning, the TRD radiators were hung onto the chambers. The radiator frame was aligned to the TRD strongback using a drill bushing fixture. During construction, the radiator frame peg holes and screw holes for mounting the frame to the aluminum strongback were drilled using a stencil. The stencil was a precision machined steel plate with drill guides, called bushings, in place to guide a hand drill. The same stencil was used to drill and tap mounting holes on the

aluminum strongback, thus ensuring proper alignment of the radiator beam holes within the overall TRD geometry.

When KTeV hall was ready, the TRDs were transported from storage and mounted on frames in KTeV hall. The TRDs were hung from rollers so that they could be rolled out during E832. When the TRDs were rolled in, their positions were fixed by a peg and hole system. This system was only good to within a millimeter, and when the TRDs were rolled back in their positions were checked externally by an optical survey team. The TRDs were cabled using flexible cable trays so that no special cabling operation was necessary during the process of rolling in and out.

4.6 The TRD Electronics

The normal operating conditions of the TRD system was to operate the anodes at about +2250 V with small plane to plane variations reflecting intrinsic gain on those planes. The cathodes were operated at +200 V with respect to ground, and the window was operated at ground. The aim of the voltage tuning was to try and maintain the pion MIP peak at about 120 ADC counts for the central gang-2 region of each front plane and at about 80 ADC counts for the central gang-2 region of each back plane. The high voltage was distributed to the anode planes through 100 M Ω resistors. The cathode planes, on the other hand, were connected to high voltage going through 667 k Ω to ground; so that during normal operation, the cathode drew 30 mA of current. This was done for the following reason. During the E799 startup in December 1996, it was found that at high beam intensities, the cathode voltages drifted *above* their nominal voltages of 200 V. While the CAEN HV supply was able to supply bias current when the cathodes or anodes were pulled under voltage, they did not have the capability to sink excess current on the cathodes during high intensity beam. This allowed the CAEN to sink excess current on the cathodes.

The TRD frontend electronics consisted of a preamp and postamp channel for each TRD channel. The preamps were capacitively coupled to the ganged anode wires at high voltage. In order to perform well, the preamp grounds needed to be grounded to

the chamber frame at as many points as possible. The output of the preamps was a differential voltage signal that traveled to the postamps through shielded twisted pair cable. The preamp had a nominal gain of 40. The signals were received on a postamp card. The postamp output was a single ended signal into RG-58 coaxial cable which fed the ADCs. The postamps also included a cable pre-compensation network designed to alter the signals in compensation for travel through the RG-58 as well as a shaper circuit designed to reduce the length of the signals as much as possible. The signals were amplified at the postamp with a nominal gain of 10. Lastly, the postamps included a pair of discriminators for each channel. The discriminator outputs were used in the TRD trigger.

While the sources of noise were many in the pit of KTeV hall, one copious source of noise were the CAEN HV supplies themselves.⁴ The noise appeared whenever the on-board SY527 computer was finished booting, and was 50-150 mV of 100-200 ns duration, every 5-10 usec. The CAEN noise was removed with the aid of low-pass filters at the anode HV input to the TRDs and with the aid of ferrite cores on the HV cables themselves.

The signals from the RG-58 were received by 1920 channels of 10 bit LeCroy 4300B FERA ADC. The gate width for charge integration was 300 ns.

4.7 CAEN High Voltage Control and Monitoring

In KTeV, high voltage to the DC and TRD systems was provided by two CAEN SY527 mainframes. The SY527 is as wide as a standard CAMAC crate and can be mounted in a standard rack. Each SY527 can hold up to 12 cards in the back which distribute high voltage. The crate itself is “smart” and contains an on-board computer to control and monitor HV distribution through a front panel LCD interface. The on-board computer also has the ability to carry on remote communications through a front panel “CAENnet” interface. The communications are serial through a standard LEMO cable. Up to 100

⁴This noise was observed only with the SY527 units, which had elaborate on-board computers. The noise was not observed on other CAEN units; for example, the NIM compatible CAEN 4 channel supply used during R&D.




Figure 29. The CAEN HV system as used in KTeV. Remote communication with the CAEN SY527 mainframes in the pit took place through a C117B CAMAC interface. An emergency kill bypass through direct front panel interface also existed but was never used.

SY527 crates can be connected with CAENnet in a daisy chain architecture where each crate in the chain is assigned a unique number from 0 to 99. Remote communication can take place through the C117B CAMAC interface. KTeV used two SY527 units in this way as depicted in figure 29. Full functionality in remote control and monitoring of the SY527s is supported by the CAENnet protocol. To do this, a CAENnet protocol “packet” is assembled in the C117B using standard CAMAC transfers and flushed through the daisy chain. The addressed crate will then respond and may assemble a response packet in the C117B. The controlling process can then poll the C117B and read the response packet using standard CAMAC transfers when the response is ready. The C117B has no command buffer and was dead while the SY527 processed a request. In principle, this is not a problem since the C117B will return an appropriate error message if it is in a busy/wait state.

4.7.1 CAEN Server Program

In the design of CAEN communications software, the main goals were flexibility and reliability. The software should allow the execution of any supported functionality of the SY527 units. Furthermore, the software should reliably execute all of the commands and report the responses. Finally, special attention was paid to certain mission critical objectives : reliability of monitoring (ie- no error codes in the returned monitoring data and reasonable response time) and ability to quickly shutdown HV in case of a problem (ie- fire, explosion, or some such disaster.)

However, the non-buffered nature of communications was a worry during CAEN control software development. I found that concurrent requests made to the C117B for CAEN access could result in data streams from the C117B containing multiple error messages. Most worrisome, I found that when the C117B was hammered with multiple requests to read out the status of channel groups, at least one of the resulting data streams would contain error codes where monitored voltages were expected. This is undesirable for two reasons. (1) Clearly, readout began when the C117B wasn't ready and the error code was not returned at the beginning of readout. Rather, the error code came in a random position along the response stream. This is indicative of loss of atomicity.⁵ (2) Even if the error condition could be ferreted out of such a corrupt stream, there is no guarantee that multiple concurrent requests could not result in the "lockout" of a mission critical request. This could happen for instance if a designated monitor process repeatedly gets error codes while some evil process is hammering the C117B with requests for trivial service.

The solution is, of course, more software. In order to avoid hammering the C117B with concurrent requests, it was considered better to buffer the requests to the CAEN and execute them one by one. This solution was implemented using a client/server architecture and inter-process communication via UNIX SystemV style message queues. The basic scenario is depicted in figure 30. The CAEN server process maintains two message queues,

⁵An *atomic* request is one that is processed from beginning to end without interruption.

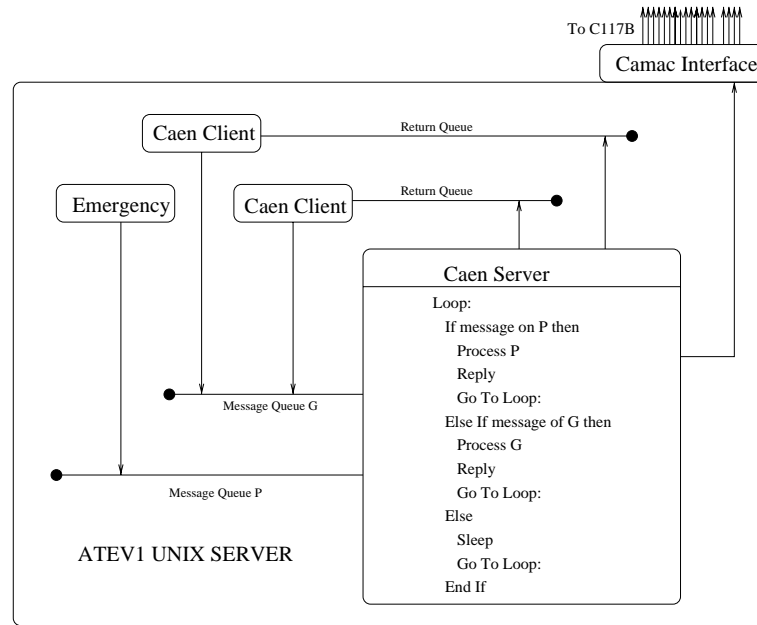


Figure 30. The CAEN server process. The CAEN server listened on open message queues. The “G” (General) queue handles run-of-the-mill requests for HV monitoring and control. The “P” (Priority) queue handles emergency requests for voltage shutdown. Since the “P” queue is checked before the “G” queue during each processing cycle, messages on “P” cannot be locked out by multiple messages on “G”. Clients supply the server with a return queue for data transfer if needed.

P and G. Messages placed on a queue are buffered and then processed in FIFO order. Response is routed back to client processes via a client provided message queue. The G queue handles general purpose requests for HV control and monitoring while the P queue handles emergency requests for HV shutdown. The processing algorithm guarantees atomicity of requests since the message queues are not checked until the current message is finished, and the atomicity of the message queues is guaranteed by the operating system. Upon finish, the server checks the P queue for a message to process. If there is a message there, the server processes it. If not, the server checks the G queue. In this way, an emergency request can bypass a string of general requests within one message processing cycle.

In order for this scheme to work, the CAEN server must be the only process running that communicates with the C117B. Since anyone can write a program using CAMAC commands, this rule was enforced by convention. There was a general purpose debugging

program that I wrote which bypassed the CAEN server. To keep people aware of what they were/might be doing with this program, I called it *cheathv*. The CAEN server program was called `caen_server.exe`. The SDAQ (described in chapter 3) main control process made sure that a copy of `caen_server.exe` was always running.

The server program had almost complete functionality. The two functions not supplied by the server were (1) channel group maintenance, and (2) killing the HV of SY527 crates. Each SY527 had the ability to maintain a list of channels in a group labeled from 1 to 12. The channels in a group list could be read out in a single transfer request, thus reducing the bandwidth required for complete readout of a detector system. Channel group maintenance, while taking place through the C117B, was however not implemented in the server for a strange reason. To add or delete a channel from a group requires up to 30 seconds.⁶ This is an inappropriate amount of deadtime. Since lists of channel groups could be defined at the beginning of the experiment and left alone forever, I decided to leave this functionality in the *cheathv* program only. Also, since the crate killing functionality was implemented directly through the front panel interface and not through the C117B, it was not handled by the server.

4.7.2 CAEN Clients

I wrote a general purpose client program called `moncaen_auto.exe` to access CAEN data through the caen server. All of the higher level DC and TRD monitoring functions were controlled using this program as the basic building block. Thus, multiple copies of `moncaen_auto.exe` can be running without concurrency conflicts.

4.7.2.1 *caen_master.tcl* GUI

The `caen_master` program provides a simple graphical user interface for general control of all aspects of the KTeV CAEN HV system. The voltage and current limiting

⁶I don't know the reason for this, but it might be that on receipt of a group maintenance command, a tiny penguin living in the SY527 has to burn a new chip and install it. 108 channels and 30 seconds per channel brings a lot of such scenarios to mind.

characteristics of each channel are controllable. Furthermore, the status of each detector system using CAEN HV in KTeV could be read out by group : the drift chamber spectrometer and the TRD system. In an emergency, `caen_master` provides the capability to kill a channel, group of channels, or a crate of channels. Also, the `caen_server` process could be killed and restarted, and the C117B CAMAC interface could be reset.

4.7.2.2 SDAQ monitoring

The SDAQ system, among other things, was responsible for routine monitoring of all CAEN high voltages for the drift chamber spectrometer and the TRD system. The SDAQ maintained a separate GUI client process to provide the shift crew with up to date information on the CAEN HV as well as automatic alarms if a channel tripped or was operating out of parameters.

4.7.2.3 TRD Alarm Process

A separate TRD alarm client process monitored TRD HV and raised an alarm if a channel was operating above or below normal operating voltages and/or currents. The TRD alarm client allowed the user to easily change the parameters defining normal operating conditions. Also, alarm monitoring of selected chambers could be easily disabled if the chamber was down or otherwise known to be operating outside of parameters temporarily. Finally, the TRD alarm process would turn a TRD chamber off automatically if alarms were not responded to in a timely manner.

4.7.2.4 TRD Active HV Control

During a long period of running, the TRDs experience gain drifts due to changes in the weather. The barometric pressure changes affects the gain because the absolute gas density changes at the anode wires. The effective form for the relation between voltage, gain, and pressure used for active HV control is

$$\Delta V = b \log(g + s\Delta p), \quad (4.1)$$

where Δp is the change in barometric pressure with respect to 745 mm of Hg, g is the gain set point in ADC counts, ΔV is the required voltage change, and b and s are constants to be determined from the data. In this model, it was assumed that the pressure changes would induce small linear changes in gain independent of voltage. The constant b was determined separately for each TRD plane in voltage scans using a small chamber mounted Fe55 source. A new voltage scan was done at the beginning of the E799 winter run and the summer run. The constant s was determined from observation of chamber gain over a period of time during which the barometric pressure changed. s was further tuned at the beginning of each run to give good performance. The values of b ranged from 150.0 to 200.0 volts, and the values of s were about 1.5 ADC counts per mm Hg.

During the run, the TRD active HV control process would read barometric pressure from a monitor in the detector hall near the TRDs. The program would guess the change in gain due to changes in pressure and compute a voltage change needed to maintain equal gain using equation 4.1. (See figure 31.) However, the active high voltage calibration drifted slightly on a much longer time scale due to small changes in gas composition. Therefore, the target gain g was fudged once or twice during the run. The performance of the active HV control kept the day to day variations in gain to under 10% of itself as shown in figure 32.

Although interesting, the active HV control of the TRD system was not important to the final offline performance of the system. It was however very important to the functioning of the level 2 TRD trigger. Large fluctuations in chamber gain could have a severe impact on the performance of the TRD trigger. The TRD trigger worked by setting thresholds on each wire, so that efficiency to pass a threshold varied with the gain. The TRD trigger was not used in this analysis.

4.8 Gas System

The active gas volumes of all eight TRDs were connected to a common reservoir of Xe/CO₂ administered by the TRD gas system. The buffer gas volumes were connected to




Figure 31. The active HV control system receives input from a barometer. During E799, the barometer was checked every 15 minutes. A gain shift was calculated based upon the shift in pressure from the last check. Using a calibrated voltage - gain curve for each TRD plane, a compensating voltage shift was calculated. The active HV process then adjusted the TRD high voltages automatically through the CAEN server.

a separate common reservoir. Both reservoirs were recirculating so that ideally, no gas was lost during the run. The gas composition was monitored daily with a gas chromatograph. The chromatograph checked for impurities which would indicate a leak and also monitored the principle components of each gas mixture. If the composition changes, makeup gas could be added to bring the gas composition back within operating parameters.

The gas system was also responsible to maintain a careful pressure difference between the active and buffer volumes at the ± 2 mTorr level. This was important to maintain the position of the active inner TRD window to within ± 50 microns during the run. The entire system was maintained at a slight overpressure so that in case a leak developed, gas would leak out of instead of in to the TRD gas system. Finally, to pressure release valves were in place on each chamber to prevent failure of gas volume integrity in case of a catastrophic accidental overpressure.

To remove impurities from the gas, online systems were employed to remove O₂, H₂O, and N₂. The O₂ was removed by commercially available systems, and the H₂O was removed by a commercially available molecular sieve. A special new N₂ remover was




Figure 32. The performance of the active HV system during the part of the E799 winter run is plotted above for TRD1. Before run 8200, the TRDs were not running reliably for various reasons. After run 8200, the TRDs were online and the active HV system was running. For a brief period around run 8480, the active HV failed. Interestingly, this active HV failure corresponds to a sharp bump in the plot.

employed for the first time in the KTeV system, due to J. Krider of Fermilab. The N₂ remover depends upon the fact that the diffusion rate of Xe through a system of silicone tubing is much faster than that of N₂. Figure 33 shows a schematic of the system. Chamber gas is piped through the tube in a sealed canister. The Xe diffuses out of the tubing, and is captured in the canister.

During normal operation, the concentrations of the nominal components of the chamber gas were maintained by injections of these components when deemed necessary. These components were tracked by the gas chromatograph as well.




Figure 33. The N₂ remover designed by J. Krider at Fermilab. The system depends upon the difference of permeation rates through silicone tubing for N₂ and Xe. Xe diffuses through more quickly and is therefore trapped in the canister. The Xe depleted gas is exhausted.

4.9 TRD Performance

After the run, the TRDs were calibrated to measure the response of the pion MIP peak from calibration $K \rightarrow \pi^\pm e^\mp \nu$ events from channel to channel. The channel by channel gains were used to equalize the channel by channel response before input into the likelihood calculation. Also at this stage, the positions of each channel were calibrated. During the likelihood calculation, downstream track segments were used to point at specific TRD channels to determine channels in which to expect a signal and to reject accidental activity.

The spectrum for pions will be characterized by a single ionization peak corresponding

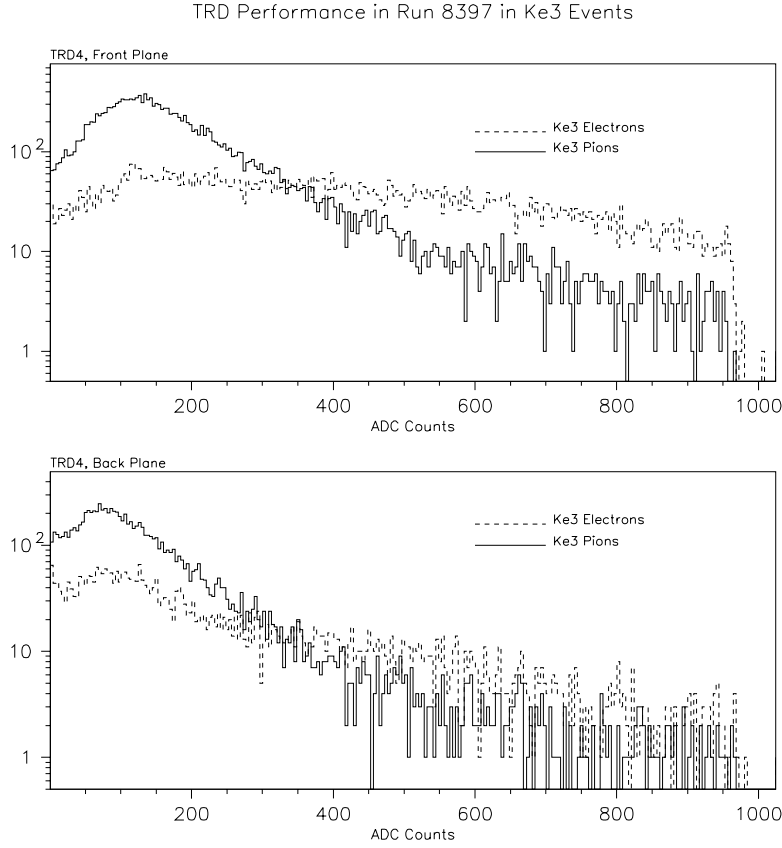


Figure 34. TRD ADC spectra are shown from the front plane (top) and the back plane (top) of TRD4. The events shown are identified Ke3 events. The electrons have a harder spectrum than the pions due to the addition of transition radiation.

to the ionization dE/dx of the pion, while the energy spectrum for electrons will be characterized by the sum of transition radiation and its ionization dE/dx . These spectra are shown in figure 34 for a representative run from E799. The figure shows the response for electrons and pions in a clean sample of Ke3 events for both the front plane and the back plane of TRD4. In order to reject pions, the responses from all 16 planes are combined into a single analysis variable called the pion probability. This method is calibrated by knowing the shape of the pion spectrum. For a given ADC measurement of plane i , x_i , the probability, P_i , is calculated that the pion spectrum could generate a signal greater or equal to x_i . For a single measurement, this is essentially a confidence




Figure 35. For two ADC measurements x_1 and x_2 , the probabilities that a pion generates a signal greater than these values are P_1 and P_2 respectively. The total probability is the product of the two individual probabilities, which generates a hyperbola in the probability space as shown above. A confidence level can be given as the integral of the probability space below the hyperbola.

level for the measurement being consistent with a pion. By construction, a sample of pions will generate a flat distribution in P_i .

The confidence levels from several planes can be combined into a single analysis variable. For a set of N measurements ($N \leq 16$), the probability of generating a set of measurements of ADC value x_i or greater is

$$P_N^{\text{tot}} = \prod_{i=1}^N P_i, \quad (4.2)$$

and the confidence level is given by

$$C_L = P_N^{\text{tot}} \times \sum_{i=1}^N \frac{-(\ln P_N^{\text{tot}})^i}{i!}. \quad (4.3)$$

In two dimensions, this can be seen as the integral of the area below the hyperbola defined by $P_2^{\text{tot}} = P_1 \times P_2$. This is illustrated in figure 35. The area below the hyperbola is the combined confidence level,

$$C_L = P_2^{\text{tot}} (1 - \ln P_2^{\text{tot}}). \quad (4.4)$$

C_L is known as the pion probability.




Figure 36. The pion probability is shown for a sample of Ke3 electrons overlaid with Ke3 pions. A cut in pion probability is used to reject pions.

From an identified sample of Ke3 events, figure 36 shows the total pion probability for a sample of pions and a sample of electrons. By construction, if all of the ADC measurements on all TRD planes are uncorrelated, then the pion probability for a sample of pions will be flat. Electrons, which tend to have higher ADC signals than pions, will have low pion probability. In practice, residual correlations between planes cause the pion sample to be slightly curled up at the edges. A cut can then be placed upon the pion probability to identify electrons. Using this method, a pion rejection of $276 \pm 38 : 1$ has been measured at 90% electron efficiency in run 8397. Figure 37 shows the rejection capability in each calibration range throughout E799. The overall single track pion rejection capability was better than 200:1 at 90% electron efficiency over the entire run.




Figure 37. The single track pion rejection performance at 90% electron efficiency is shown in all of the E799 KTeV TRD calibration ranges. The winter run comprises calibration ranges 08a through 33b. The summer run comprises calibration ranges 37b through 55a. *Plot courtesy of L. Bellantoni.*

CHAPTER 5

Elementary Event Reconstruction

In this chapter I will describe the elements of event reconstruction. Since proper calibration is essential to reliable event reconstruction, I also describe briefly what needs to be calibrated and how those calibrations are done. Later in chapter 7, we will investigate the effects of shifts in calibration with respect to systematic error.

5.1 The Database

A database based on the CERNLIB HEPDB system was used to store all of the online and offline calibration constants in KTeV. A common format was imposed on all of the calibration constants. It is also at the center of the capability of KTeV analysis and monte carlo software to reliably obtain the correct set of calibration constants for any given period in the run. A set of calibration constants in this scheme is keyed by detector, run number, and isource. Each bank of constants has a range of valid runs, used to key off of the run number. Isource = 0,1 selects the “online” constants which were generated during the run and used in the L3 filter. Isource=2 selects the better “offline” constants generated later. A constraint is also placed on the selection of calibration constants by date; implicitly, the database software will try to return the most recent calibration constants consistent with the given keys. One can specify a date beyond which the database will not go to retrieve constants.

This architecture makes it particularly easy to simulate different aspects of the analy-

sis. For example during monte carlo generation, the L3 filter is simulated using the online constants. Also, the crunch (chapter 3) is simulated by specifying the actual date of the crunch. Since the real crunch did not have access to calibration constants generated after it ran, the post-crunch constants are not used in simulating the crunch.

During analysis or monte carlo simulation, a run number is either read from the data or generated (see chapter 6.) If the run number has changed, a routine is called to check if the current constants are still valid. Any constants for which the current run number is outside of the validity range are reloaded from the database.¹

Some of the more basic calibrations will be discussed below in conjunction with the event reconstruction.

5.2 Event Reconstruction

A typical KTeV analysis can reconstruct entire kaon decays. This primarily involves the finding of charged tracks in the spectrometer and energy clusters in the CsI. A sample algorithm appears below followed by brief explanations of the reconstruction routines.

5.2.1 A Sample KTEVANA Job

```

SUBROUTINE KTANA

C
C  General Analysis : Tracking and Clustering
C    An event has been loaded.  Calibration constants depending
C      on the run number have also been loaded.
C
C  Tracking I : T3TRAK
C
NWANT = 0

```

¹There were two notable exceptions to this rule : Trigger definitions and certain beamline constants were read in by “proprietary” files and did not reside in the HEPDB database.

88

```
      LSHAREY = .true.
C
C   Tracking analysis
C
      Call T3TRAK(NWANT,LSHAREY,TRAKSTAT)
C
      If (TRAKSTST.ne.0) Go To 999
C
C   Find clusters ...
C
      IPACK(1) = 16447
      IPACK(2) = 31
      NSLICE = 4
      CLUSTER_THR = 0.25
      BLOCK_THR = 0.1
C
      Call KTCLUS(ipack,NSLICE,CLUSTER_THR,BLOCK_THR,0,0,IERR)
      If (IERR.ne.0) Go To 999
      Call Record_History('KTCLUS IERR = 0')
C
C   Find vertices
C
      Call T3QVTX(LSHAREY,QVTXSTAT)           ! find vertex candidates
      If (QVTXSTAT.ne.0) Go To 999
C
      MINMCODE = 3                             ! require tracks
                                           ! to match clusters
                                           ! with none in beam
                                           ! hole at CsI
      Call T3FVTX(MINMCODE,LSHAREY,FVTXSTAT) ! find best vertex
```

```

! and corrected tracks

      If (FVTXSTAT.ne.0) Go To 999
C
C      Your Analysis Here !!!
C      Call 773-702-5983
C      Ask for V. Prasad
C
999  Continue      ! Event failed reconstruction.

      Return
      End

C

```

After the above general purpose analysis requirements, the number of tracks and clusters and their positions, momenta, and energies are available for individual analyses.

5.2.2 Clustering

5.2.2.1 Energy Calibration

The clustering algorithm uses 4 DPMT time slices of readout to determine cluster energies. The energy calibration done for these analyses took this into account by also using four slices of information. Typically, four slices contained about 95% of the energy in a cluster. A 6 slice calibration was also considered, but the improvement in energy resolution was negligible with respect to this analysis.

A calibration was done to convert DPMT counts into phototube charge. This calibration was done using the laser system. During the run, many scans of the array were done with the laser to obtain this calibration.

A second set of constants is needed to convert charge into energy. This calibration used electrons from K_{e3} events. The general procedure was to calculate the momentum




Figure 38. The energy calibration in a selected run range of E799. The calibration uses the property that electro-magnetic showers are completely contained in the CsI, so that the mean ratio of energy to momentum, or E/p , for electrons is 1.0. The figure at left shows the E/p distribution for electrons from K_{e3} decays before the offline energy calibration. The figure at right shows the distribution after offline calibration for the same set of data. The mean E/p after calibration is 1.0006 by Gaussian fit, and the resolution is 0.99%.

of electrons using the charged spectrometer.² These electrons made clusters in the CsI. The cluster energy was estimated using a rough set of calibration constants and E/p was calculated. The deviation of E/p from 1.000 is a measure of the miscalibration of the crystals that form the cluster. Over a calibration run, these differences are averaged for each crystal and the averages are used to produce an improved set of calibration constants. The energy calibration is depicted in figure 38.

5.2.2.2 Algorithm

The clustering algorithm used in these analyses looks for *hardware* clusters, or clusters in which the HCC bit was set. A list containing all of the crystals above HCC threshold is searched for local maxima. A local maximum is a crystal with energy deposit greater than any other crystals in the list that share an edge with that crystal. The crystals which are identified as hardware local maxima were called cluster seeds.

²Interestingly, this momentum calibration depends on a calibration of the Δp momentum kick of the analysis magnet, which comes from $K_L \rightarrow \pi^+ \pi^-$.

The raw cluster energy is the sum of energies in a square region around each cluster seed. For small crystals, the region was a 7×7 square array; for large crystals, it was a 3×3 square array. The ratios of seed block energy to neighbor block energies was used to calculate the cluster position. The positions were taken from lookup tables, and the resolution was about 1 mm.

5.2.2.3 *Energy Corrections*

There were several corrections applied to the data in these analyses.

- **Overlap Correction** : Sometimes, a cluster may share crystals with a neighboring cluster. This correction estimates and corrects for the energy shared between neighboring clusters. The estimate is based upon the expected electro-magnetic shower shapes from GEANT.
- **Neighbor Correction** : Sometimes, a particularly large energy cluster may deposit energy in crystals outside of the standard 3×3 or 7×7 array of crystals used in the clustering algorithm to contain it. If a smaller non-overlapping crystal is nearby, its energy may be significantly enhanced by this effect. Therefore, a correction is also applied to the energy of the smaller cluster to account for this effect. This correction is also based upon the expected shower shapes using GEANT simulation.
- **Intrablock Correction** : The energy of a cluster was scaled to account for small deviations in cluster energy observed when the cluster position was near a crystal boundary as opposed to the center of a crystal. This correction is taken from effects observed in the data.
- **Threshold Correction** : As described in chapter 3, the digital pipeline system imposed a readout threshold of 12 counts to the CsI data. This corresponds to about 7 MeV of energy. Clusters that include crystals below threshold are corrected for this effect. The correction was determined by taking data in a special run with no readout threshold in the digital pipeline. The energy of under-threshold crystals




Figure 39. The threshold correction is depicted for a small crystal that is located 3 blocks above and three blocks to the side of the seed block. The mean energy as a fraction of the applied threshold is plotted versus cluster energy for under threshold blocks. The data shown are electrons from K_{e3} events in run 9897, a special run with no CsI readout threshold.

was then fit as a function of cluster energy and crystal position in the cluster. The crystal position in this correction is indexed by number of crystals above and below the seed block. (For large crystals, where a cluster was found in 3×3 arrays, crystals were almost never below threshold.) To correct the cluster energy, the expected energy from under threshold blocks is added back in.

- **Missing Block :** Sometimes a cluster is missing one or more blocks. This happens when the 3×3 or 7×7 array of crystals used by the algorithm to contain a cluster is located near an edge of the calorimeter. This correction adds the energy back into the cluster that would have been deposited in the cluster based on expected shower shapes generated in GEANT.
- **Run by Run Correction :** The quiescent laser data was capable of tracking variations in crystal response with respect to time. The variation was most evident in the crystals in the center of the array which received the brunt of the radiation dose from kaon decays. The quiescent laser corrections were not performed for E799, however, so the main method to combat this problem in E799 was to calibrate often. The short calibration run ranges were chosen so that the mean E/p response of the central crystals varied by no more than 1.0%. Also, a direct correction based

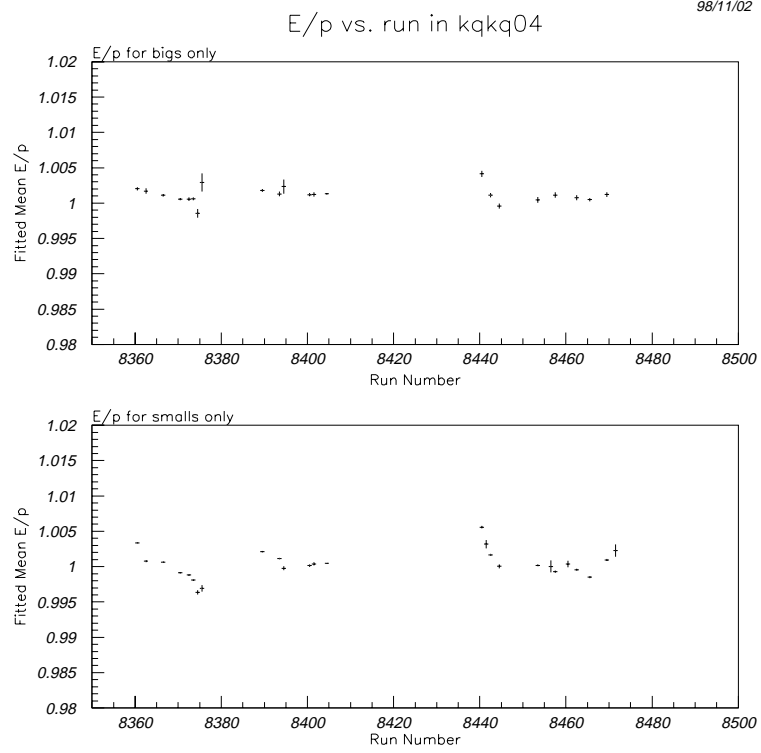


Figure 40. Mean E/p versus run number for E799 runs in the “kqkq04” range, an arbitrary run range. The top plot shows E/p in clusters consisting entirely of large crystals and the bottom shows E/p in clusters consisting entirely of small crystals. (The relevant calibration ranges are (a) 8358 to 8377, (b) 8378 to 8407, (c) 8408 to 8451, and (d) 8452 to 8482.) The run to run variations are much greater in the small crystal region in the center of the CsI array where radiation damage is greater. The calibration ranges were chosen to limit the overall variation to 1% in any range, and corrections were later applied based on plots like these within each range.




Figure 41. Mean E/p is plotted versus electron momentum for electrons from K_{e3} events in clusters consisting of large crystals only. This is the basis for the final non-linearity fudge for the large crystals.

upon the observed variation within a calibration run range was applied. The regions in the CsI array over which the correction was applied were (a) the region where clusters consisted of large crystals only, and (b) the region where clusters consisted of small crystals only. This correction was calibrated and applied in this analysis only and was used in systematic studies of the energy scale.

- **Non-Linearity Fudge :** After all of the standard corrections were applied, there remained a small non-linear dependence of E/p on p . This effect was measured for the different calibration run ranges. The typical variation was on the order of 0.5%. This residual measured non-linearity was used to linearize the E/p response. The regions in the CsI array over which the correction was applied were (a) the region where clusters consisted of large crystals only, (b) the region where clusters




Figure 42. Mean E/p is plotted versus electron momentum for electrons from K_{e3} events in clusters consisting of small crystals only. This is the basis for the final non-linearity fudge for the small crystals.

consisted of small crystals only, and (c) the region where clusters consisted of both large and small crystals. The correction was calibrated and applied in this analysis only and was used in systematic studies of the energy scale.

5.2.2.4 Conclusions

The final energy calibrations feature an E/p peak whose mean is within 0.1% of 1.000, (see figure 38,) for all calibration run ranges. Furthermore, the E/p width is about 1.0% for all run ranges.




Figure 43. Good SOD, Low SOD, High SOD, No SOD. The different kinds of SODs are shown. Typically, a single track traverses a pair of drift cells so that the sum of drift distances add up to $\frac{1}{2}$ of a cell spacing. However, if two tracks or a single track with a delta ray are present, a low SOD pair will result. The high SOD pair occurs when a track passes very close to a wire as explained in the text. Finally, if a wire is inefficient, an isolated hit can result.

5.2.3 Tracking and Vertexing

5.2.3.1 Hits and SODs

The tracking algorithm starts by looking for hits in the drift chamber TDCs. Although each TDC channel had the capability of storing multiple hits, only the earliest hits in the in-time window (see figure 11) were used by the tracking algorithm. The valid hits were then converted to drift distances using the XT maps. The XT maps were a set of calibration constants constructed for this purpose, and are discussed briefly below.

When a track traversed a pair of drift planes it usually left a pair of hits, one in each plane of the view, (see figure 43.) From the geometry of the drift cells, the sum of distances (or SOD) should be equal to $\frac{1}{2}$ of the cell spacing, or 6.35 mm. A hit pair was classified according to its SOD. If the SOD was close to the nominal 6.35 mm, the hit pair was considered a “good SOD” pair. Sometimes, two tracks could traverse the same cell. In this case, each sense wire would see the charge from the nearest track and the SOD would be less than the nominal value, and the hit pair was classified as a “low SOD” pair. Low SOD pairs could also be caused by delta ray emission. Finally, there were also “high SOD” pairs. High SOD pairs could arise when a track passed close to a wire in one of the complementary planes. If it was close enough to the wire that the mean




Figure 44. A comparison of SOD distributions between E799/E832 and the previous experiment E773. Both experiments used the same chambers and wire geometry.

spacing between ion pairs created by the track in the drift gas was non-negligible, a high SOD could result. This effect has been modeled in the monte carlo. Furthermore, the ionization potential of the DC gas has been fudged in bulk regions of the chamber to make the SOD distributions agree between data and monte carlo. The rationale for calibrating in different DC regions was to take full account of the varying rate effects in different DC regions. This calibration has been successful at modeling most of the high SOD tail, and was done for selected runs in E799. However, it does not model the very-high (greater than 2 mm) high SODs very well. It does do well enough for the analyses presented here.

Figure 44 shows that the high SOD problem is much worse in E799 than in E773. In E799, about 3% of the hit pairs have SODs more than 1 mm high. In E773, the figure is less than 0.5%. The cause of the “very high” SOD problem has not been conclusively determined, and it has not been fully modeled in the monte carlo.

A supplementary method has been proposed which empirically creates maps of high

SOD “probability” in each drift plane. The maps essentially compute the probability that a track will produce a high SOD pair in 1×1 cm bins across the face of each drift plane, and duplicates this behavior in the monte carlo. The high SOD maps have been implemented in E832, and the work is in progress for E799. Explanations for the high SOD problem invoke all of the usual suspects as radiation damage or voltage sag in the beam region.

In addition to the hit pairs, sometimes a hit is missed in one of the two planes of a plane pair due to intrinsic inefficiencies. In this case, the hit is classified as an isolated hit. Because isolated hits carry ambiguous position information, a track can be constructed from some isolated hits if good SOD pairs are also present but clearly cannot be constructed from only isolated hits. Lastly, if a pair of track candidates points to a low SOD pair, the hit pair will be treated as a pair of isolated hits.

5.2.3.2 XT Maps and Alignment

The XT maps were calibrated using muons from special muon runs. (The muon runs are data taken with the beam on but the beam stop closed.) These are calculated by examining the distribution of hit times in each drift chamber cell. It is assumed that the distribution of hit positions within each cell is flat, and then the XT relation can be reasoned out from the distribution of hit times using bins of equal area.

Once the XT maps are calculated, the muon runs can also be used to map out the positions and alignment of the drift chambers. With the analysis magnet turned off, the straight line segments made by muons were used to find the X and Y positions of all of the drift cells as well as any rotations and non-orthogonalities. The Z positions of the chambers were found using external optical survey techniques.

5.2.3.3 Y and X Tracks

After the list of hit pairs has been tabulated, X and Y track candidates are constructed. The Y tracks are calculated first since there are no bends in the Y view. Hits that were

collinear within 0.5 cm between DC1 and DC4 were considered for Y track candidates. The hits were allowed to include at most two low SOD pairs or at most one high SOD pair and one isolated hit.

In the X view, track candidates were found by making segments between DC1,DC2 and between DC3,DC4. Pairs of upstream and downstream segments were considered an X track candidate if they intersected within 0.6 cm at the magnet bend plane. An early requirement on the upstream segments (to reduce the number of segments considered) was that the angle of the segment with respect to Z was required to be less than 100 mRad. The downstream segments were required to make an angle of less than 150 mRad with respect to the Z axis.

At this stage, there is no way to associate particular Y track candidates with specific X track candidates. Only after the CsI cluster positions are known can the ambiguities be resolved. Also, by option, X tracks were allowed to share a single Y track. Such a configuration is possible if the plane containing two tracks is oriented so that the Y tracks coincide. In the L3 and crunch analyses, Y tracks were allowed to be shared.

5.2.3.4 *PTKICK calibration*

The exact momentum kick at the magnet midplane was calculated on a run by run basis using $K_L \rightarrow \pi^+\pi^-$ events. The invariant mass of the 2π event is calculated, and Δp_{mag} is chosen so that the mean of the invariant mass spectrum is equal to the kaon mass. Since the invariant mass essentially depends only on momenta measured by the spectrometer, this calibration is very clean. The momentum kick is calculated on a run by run basis, and was typically 205 MeV/c.

5.2.3.5 *Vertex Finding*

The intersection between each pair of track candidates in X and Y were calculated to find the vertex candidates. A vertex candidate consists of pairs of X and Y view track candidate pairs whose intersections are consistent with each other and in the decay region.

Vertex candidates were also required to have opposite-bending tracks in the X view (opposite charged tracks) in the L3 filter.

After the vertex candidates were identified, those which matched cluster positions in the CsI were kept and the rest rejected. Each track extrapolated to the CsI was expected to match a cluster to within 7 cm. For the vertex candidates passing matching, the tracks were corrected as described in the next subsection.

Using the corrected track information, a χ^2 was calculated based on the calculation of the Z position of the vertex, the matching of upstream/downstream X view tracks at the magnet, and the number of good SOD, low SOD, or high SOD pairs used in the vertex candidate. The candidate with the best χ^2 was taken as the final vertex and set of tracks.

5.2.3.6 Tracking Corrections

Several corrections were made to the hit positions. The hit positions in DC2 and DC3 were adjusted to account for fringe fields from the analysis magnet between DC1,2 and between DC3,4. The fringe fields were mostly in the vertical direction and produced a kick of about 1 MeV/c. Hit positions were also corrected for chamber alignment and for corkscrew rotations.

CHAPTER 6

Monte Carlo Simulation

The KTeV monte carlo simulation is used to find the acceptance of kaon decay modes and to study various systematic effects. The acceptance for a given decay mode is the efficiency for reconstructing the decay given the detector geometry and the analysis criteria. The calculation of the sensitivity for the $K_L \rightarrow \pi^0 e^+ e^- \gamma$ decays will involve calculating the ratio of acceptances of the signal mode to its normalization mode, $K_L \rightarrow 2\pi_D^0$. The acceptance for various background processes can be estimated using the monte carlo, and measured decay distributions can be checked against the monte carlo decay distributions to look for and to estimate systematic error.

It is interesting to think of the acceptance as the integral of a function, F , of the detector geometry, G , and analysis requirements, R . For a given event e , we say that

$$F(e; G, R) = \begin{cases} 1 & \text{if } e \text{ passed requirements } G \text{ and } R \\ 0 & \text{if } e \text{ was rejected} \end{cases} \quad (6.1)$$

In order to measure the overall acceptance for a decay mode, the acceptance function F is integrated over the domain of kaon decays. The domain of kaon decays is defined by an input momentum spectrum, K_L - K_S transport to the detector, and a set of generated four-momentum vectors of decay products in the lab frame. To determine the value of F for a given event e , the products of a particle decay are traced through the KTeV detector and detector response is simulated. The event is accepted if it passes all of the trigger and analysis requirements. We can estimate the integral of F over the set of generated events E using the Monte Carlo method of integral estimation. The acceptance is then

estimated by

$$I = \frac{\sum_{e \in E} F(e; G, R)}{\sum_{e \in E} 1}, \quad (6.2)$$

where the denominator is just the number of generated events in the domain of integration. The standard error on this number goes as \sqrt{N} where N is the number of accepted events in the numerator. However, the error estimates become a little more complicated if F is allowed to be a weighting function that takes on values between 0 and 1. (See appendix A.) This will be the case for a large part of the $K_L \rightarrow e^+e^-\gamma\gamma$ monte carlo generated for this analysis, which is the primary background for $K_L \rightarrow \pi^0 e^+e^-$ and which was considered as a background to $K_L \rightarrow \pi^0 e^+e^-\gamma$.

In this chapter, we will discuss the general aspects of the KTeV monte carlo first and then delve into the individual decay generators. To get a realistic estimation of F , it is necessary to model the detector in great detail.

6.1 General Aspects of the Monte Carlo

In this section, the algorithm for generating a monte carlo event is described. Some of the more important aspects of this, including detector simulations, are described in more detail afterwards. The generation of monte carlo events is, in a way, the inverse of the analysis process presented in the last chapter. There, we were concerned with the use of calibration constants to turn measured detector responses into physical quantities. Here, we are concerned with the use of the same calibration constants to generate fake detector responses from calculated physical quantities.

6.1.1 The Life and Times of a Monte Carlo Event

The monte carlo is initialized with a choice of decaying particle, decay mode, momentum range, and Z range. The momentum range tells the monte carlo which range of initial kaon momentum is desired; and in these analyses, it is always 20 GeV/c to 220 GeV/c. Below 20 GeV/c, K_L tend to decay upstream of the detector and the decay products tend to escape the detector and/or fire the photon vetoes. The Z range tells the monte

carlo where the decay region is; and in these analyses, it is always from 90.0 m from the target to 160.0 meters from the target, (see figure 9). The decay mode specification will tell the monte carlo which subroutine to call to generate decay products and their lab frame 4-momenta. The particle specification tells the monte carlo how to select the initial decaying particle.

The analyses presented here use the “K0K0BAR” mix for the decaying particle. With this choice, the monte carlo program will choose between two possible particles to generate for each event, K^0 or \bar{K}^0 . This designation makes physical sense since these are the production states at the target and have different momentum spectra. The distinction in the monte carlo program is twofold. The choice of K^0 or \bar{K}^0 determines the initial vector in the K_1 - K_2 basis used in the kaon amplitude evolution, and it determines the set of fit parameters to use for the in momentum and angle distributions. The relative weights for making a K^0 or \bar{K}^0 are fixed by evaluating the probability that a generated K^0 or \bar{K}^0 will be produced at the target and decay in the fiducial volume defined by the above momentum and Z ranges assuming full kaon amplitude evolution. The amplitudes are column matrices in the K_1 - K_2 basis, so that K^0 is represented by $(1, 1)^T$ and \bar{K}^0 is represented by $(1, -1)^T$ ignoring normalization. Once the K^0 and \bar{K}^0 relative weights are chosen, a 2D histogram of probability to generate a decay kaon with given momentum and direction at the target is generated. These histograms are later used as random lookup tables to quickly pick an initial momentum and direction vector for the kaon at target.

Other input includes the name of a file (the MRN file) that contains pointers to files of accidental events. These accidental files contain data events taken in the accidental trigger during E799 indexed by run number. In the monte carlo, these events are overlaid on generated events to simulate the effects of accidentals.

After initialization, an event loop is entered. The loop exits when the requested number of events has been accepted. At the top of the loop, an accidental event is read in from one of the accidental files. The run and spill number of the generated event is set by the run and spill number of the accidental event. The accidental files are sampled by relative weights specified in the MRN file, which tried to mimic the actual distribution of events across run numbers in E799. If the number of events accepted in a given run

exceeds the corresponding weight in the MRN file, the next accidental file is opened. In this rather complicated way, the generation run number is incremented. This is an important point, however. If any calibration constants need to be updated to reflect a new run and spill number, then it is done here.

Next, an event is generated. The decaying particle, K^0 or \bar{K}^0 , is picked according to its weight. A beam hole, $+x$ or $-x$, is chosen at random. The two beams have equal weight. Finally, the kaon momentum and direction is picked from the distributions determined during initialization. It is checked that the generated kaon will pass through the collimators up to the beginning of the specified Z range. If it does not, then the event is thrown out and generation begins anew. The simulation includes scattering in the Pb absorber.

By option, the simulation also includes the possibility that a kaon will scatter in the defining collimator. The possibility of K_L regeneration into K_S is also considered. If the option is enabled, then the scattering angle and kaon amplitude vector will be changed to reflect these possibilities. In the analyses presented in this thesis, collimator scatters are not a serious problem; therefore for most of the monte carlo I generated, this option was off. It was important, however, to verify this claim with special monte carlo featuring *only* collimator scatters. This is discussed further in chapter 7.

The generated kaon was transported up to the beginning of the decay region using the full amplitude evolution of the K^0 - \bar{K}^0 system. The simulation evaluates 2×2 transport matrices which include oscillation of the K^0 - \bar{K}^0 system and K_S regeneration in any matter present.

If the kaon makes it to the beginning of the specified Z decay region, then the exact decay vertex is determined using full amplitude evolution for $K_L \rightarrow 2\pi_D^0$ and pure K_L exponential for $K_L \rightarrow \pi^0 e^+ e^-$ and $K_L \rightarrow \pi^0 e^+ e^- \gamma$. A decay generator is then called to determine a set of initial four momenta for the decay products. Also at this point, the number NGEN is incremented for use in the denominator of equation 6.2. We then stop tracing the kaon after the decay point and start tracing its decay products through the detector. The detector responses are simulated and the effects of accidentals are added in. The trigger is evaluated, but only after the addition of accidental energies. Also,

due to the long time involved in simulating the CsI fully, the trigger is first evaluated without the E_t requirement or the HA veto requirement. If the event fails this subset of the trigger, it is rejected. If the event passes, only then the CsI is simulated and the rest of the trigger is evaluated. If the event passes L2, it is accepted. The L3 filter is then run on the accepted events, and will apply the appropriate L3 tags to the output.

6.1.2 The Momentum Distribution and K0K0BAR Mix

The momentum distribution and K0K0BAR mix is obtained using the Malensek parameterization of charged kaon production using 450 GeV proton on Be. The data are found in [Aherton 80], and the fit in [Malensek 81]. The cross sections for production of neutral kaons can be related to those for charged kaons by simple quark counting arguments. A critical assumption behind this argument is that in the case of forward scattering, the kaon production takes place in the incident proton and not in the target nucleon, (see appendix B of [Briere 95].) This assumption is good at our energy of 800 GeV/c protons on target.¹

The quark counting argument proceeds as follows. We assume that a u or d quark can come from the valence or the sea, and that the s quark and all anti-quarks must come only from the sea. Using isospin invariance, we assume that the $u\bar{u}$ and $d\bar{d}$ content in the sea are equal. With weights for individual quarks in the produced kaons defined as in figure 46, we can relate the production cross sections for neutral and charged kaons.

$$K^0 \sim \frac{1}{2}(K^+ + K^-), \quad (6.3)$$

$$\bar{K}^0 \sim K^-. \quad (6.4)$$

The momentum distributions and angular dependencies are then estimated using these relations. The Malensek parameterization for charged kaon production at momentum p and at polar angle θ is given by [Malensek 81]

¹In any case, the argument leads to expressions that describe the data well. NA31, an experiment that used target generated K_S , depends more strongly on the exact K^0 - \bar{K}^0 content of the beam in the measurement of the K_L - K_S interference term. Their neutral kaon distributions were well fit from charged kaon distributions using the following argument.




Figure 45. The Malensek production cross sections used to pick initial momentum and direction of generated kaons in the KTeV Monte Carlo simulation for the “K0K0BAR” option. The distributions also reflect weighting from the probability of decay in the specified momentum ranges and Z ranges. Thus, the survival of high momentum K_S into our decay volume is represented by small enhancements in both the K^0 and \bar{K}^0 plots.

Incident Proton		Relative Weight
Valence Quarks	u	1
	u	1
	d	1
Sea Quarks	$u\bar{u}$	x
	$d\bar{d}$	x
	$s\bar{s}$	y

(a)

Kaon	Quarks	Weight
K^+	$u\bar{s}$	$(2+x)y$
K^-	$\bar{u}s$	xy
K^0	$d\bar{s}$	$(1+x)y$
\bar{K}^0	$\bar{d}s$	xy

(b)

Figure 46. A relationship between the neutral kaon production cross sections and the measured charged kaon cross sections can be estimated by simple quark counting arguments in the beam proton. The beam proton has two u quarks and one d quark in the valence. The sea quark content relative to the valence quarks can be parameterized by x for the $u\bar{u}$ or $d\bar{d}$, and y for the $s\bar{s}$, as shown in part (a) of the figure. Using the quark assignments for the charged and neutral kaons, one can deduce the relationships shown in part (b) of the figure and those in equation 6.4.

	K^+	K^-
A	2.924	6.107
B	14.15	12.33
M^2	1.164	1.098
D	19.89	17.78

Table 13. The parameters found by Malensek to fit the production of charged kaons by protons on a Be target. See equation 6.5.

$$\frac{d^2 N}{dp d\theta} = \frac{B}{800\pi} x \frac{(1-x)^A (1+5e^{-Dx})}{(1+\frac{p_t^2}{M^2})^4}, \quad (6.5)$$

where x is the ratio of p to incident beam energy, p_t is momentum of the produced kaon transverse to the beam direction. The other parameters are fit from the data in [Aherton 80], and appear in table 13. The expressions used for neutral kaons is in analogy to equations 6.4,

$$\frac{d^2 N_{K^0}}{dp d\theta} = \frac{1}{2} \left(\frac{d^2 N_{K^+}}{dp d\theta} + \frac{d^2 N_{K^-}}{dp d\theta} \right), \quad (6.6)$$

$$\frac{d^2 N_{\bar{K}^0}}{dp d\theta} = \frac{d^2 N_{K^-}}{dp d\theta}. \quad (6.7)$$

These dependencies were also tuned using KTeV data in $K_L \rightarrow \pi^+ \pi^-$. A correction was made both to the momentum dependence as well as to the theta dependence. The momentum correction is a simple fourth order correction given by

$$\chi_p(p) = 1.0 + b_1(p - 70.0) + b_2(p - 70.0)^2 + b_3(p - 70.0)^3 + b_4(p - 70.0)^4, \quad (6.8)$$

where p is the momentum in GeV/c and the b parameters were fit from the data. (Depending on whether p was greater or less than 70 GeV/c, a different fit was used.) All of the correction parameters can be found in [KTeVana], and will not be listed here. The size of the correction was typically 5-10%. The theta correction was a simpler first order correction given by

$$\chi_\theta(\Delta\theta) = 1.0 + 0.0388\Delta\theta, \quad (6.9)$$

where $\Delta\theta$ is the angular deviation in mRad from the nominal targeting angle of 4.8 mRad.

The momentum distribution of generated kaons is shown in figure 8. The size of the theta correction was typically 1-2 %.

Once the cross sections are known, the relative mix of K^0 and \bar{K}^0 for the K0K0BAR monte carlo was obtained by integrating them over the probability of decay in the stated decay volume. The mixture is roughly 55% K^0 and 45% \bar{K}^0 .

6.1.3 Particle Tracing and Detector Response

6.1.4 Z Position of Decay Vertex

The Z position of the decay vertex is generated according to the full evolution of the K0K0BAR amplitudes at the given momentum. For decay modes in which K_L - K_S interference is important, the full amplitude evolution is extended into the decay volume. This set of decay modes includes those which are accessible to both K_L and K_S with comparable probability. For example, the $K_L \rightarrow 2\pi_D^0$ normalization mode falls into this category since $\pi^0\pi^0$ decay accounts for about 30% of all K_S decay modes. For kaon decay modes in which K_S contamination is small, a pure K_L exponential decay distribution is assumed. Since the Z distribution depends upon the decay being studied, an explanation of the generated Z distributions will be discussed below.

6.1.5 Particle Tracing

Once the decay products are enumerated along with their 4-momenta in the lab frame, the decay products are traced through the detector. Tracing of a decay product stops if (a) the particle is stopped in a detector, (b) the particle is decayed or converted in dead detector material, or (c) the particle escapes the detector altogether.

During tracing, charged particles can undergo multiple scattering. This was simulated using the Moliere theory, reviewed in [Scott 63]. Multiple scattering is simulated in bulk regions of the detector described by a radiation length and is done separately at wire planes. Wire planes are simulated with a single probability for hitting a wire, (ie - the wires are not localized).

Material simulated	Z position (m)	Total Rad. thickness (X_0)
Vacuum window	158.89	0.00156
“Air gap”	159.25	0.00147
Drift Chamber 1	159.40	0.00050
	159.48	0.00045
Helium Bag 2	159.48 to 165.55	0.00162
Drift Chamber 2	165.55	0.00045
	165.63	0.00045
Helium Bag 3 before magnet	165.63 to 170.01	0.00119
Helium Bag 3 after magnet	170.01 to 173.99	0.00119
Drift Chamber 3	174.65	0.00090
Helium Bag 4	174.65 to 180.02	0.00156
Drift Chamber 4	180.55	0.00090

Table 14. Radiation lengths of selected components following the KTeV decay volume. Radiation lengths are modeled as planes that contain the compressed amount of material from the surrounding regions. The helium bags are modeled with four such planes to better model the continuous nature of the physical regions. *Courtesy of E.D. Zimmerman.*

For electrons, bremsstrahlung was simulated using the Bethe-Heitler cross section for the energy of the emitted photon. The angle of emittance for bremsstrahlung is small, and can be approximated by

$$\frac{d\sigma}{d\theta} \sim \theta^{-4}. \quad (6.10)$$

In the KTeV monte carlo this angle was not simulated, and the photon was considered collinear with the parent electron. The Bremsstrahlung probability is described by the radiation length parameter. Radiation lengths in the detector were simulated at discrete planes which contained the “compressed” radiation length of the surrounding regions. These are listed in table 14.

The radiation lengths of regions immediately downstream of the vacuum window were measured by a novel technique due to E.D. Zimmerman [Zimmerman 99].

For photons, conversion to an e^+e^- pair was simulated using the Bethe-Heitler re-

lation. The conversion probability was based on a scaled radiation length, the usual $X_{conv} = X_{rad}$ was used [Leo 87]. angle of emittance of the electron-positron pair was simulated using code ported from the EGS4 package [Bielajew 91].

Upon the destruction of a particle during tracing, the particle was stopped and replaced by its decay or conversion products. These products were subsequently traced through the detector. When a particle is stopped in a detector element, its response is simulated and the result is digitized. If the detector was in the trigger, the threshold is simulated and a trigger latch bit set if the response was above threshold.

6.1.6 Photon Vetoes

Electrons or photons which hit the photon vetoes; RC6-10, SA2-4, and CIA; were stopped there. The response of the detector in MIPs to a photon or electron of given incident energy was calibrated in [Hanagaki 98] using $K_L \rightarrow \pi^+ \pi^- \pi^0$ events with a missing photon. The detector response was smeared by a Gaussian distribution whose width was also measured in that data. Other charged particles were simulated as single MIPs. The detector response was determined using calibration constants derived from special “muon” runs in which the KTeV beam stop was rolled in during running to stop the hadronic component of the beam. The detector response is again smeared using a Gaussian width determined in the data.

6.1.7 Drift Chambers

When a charged particle traversed a plane of drift chamber wires, the drift distance to the closest wire is converted to a drift time using the XT maps. The result is then smeared by a Gaussian based on the measured position resolution of the drift plane. To simulate inefficiency, a hit on a given wire masks out all subsequent hits on the same wire for 235 ns. Delta rays were also simulated. These could produce spur tracks that ionize gas closer to a wire than the original track and thus produce low-sod pairs. Delta rays produced in one cell were not allowed to propagate to other cells.

The magnet was simulated by producing a constant kick of about 205 MeV/c to charged tracks at the magnet midplane, $Z=170\text{m}$.

Once the list of drift chamber hits was compiled, the level 1 trigger DCOR simulation was run. If the trigger requirements were satisfied, the appropriate trigger latch bits were turned on.

6.1.8 Hodoscopes

Charged particles traversing scintillator paddles in V, V', or the muon counters were simulated by turning on the appropriate trigger latch bits. Each counter plane had an overall efficiency to record a hit, based upon the efficiencies as measured in the data. The HA hodoscope is discussed below in connection with the CsI.

6.1.8.1 TRDS

The TRDs were simulated as $0.14X_{rad}$ total material in eight planes between DC4 and the VV' hodoscopes. In the beam region where the radiator material was cleared away, the radiation length was $0.04X_{rad}$. Scattering off of TRD wires was not simulated. The electron tagging capability of the TRDs was not simulated; rather, this performance was measured in the data.

6.1.9 CsI Calorimeter

When a particle struck the CsI calorimeter, several different things could happen based on the species of the incident particle. Photons and electrons generated electro-magnetic showers. Muons would leave a MIP signal; pions could also leave a MIP signal but also could initiate a hadronic shower.

6.1.9.1 *Electro-magnetic Showers*

The electro-magnetic showers were modeled using a library of GEANT generated shower shapes. The difference between electrons and photons is characterized by the different depths which they penetrate the CsI before initiating a shower. The photons travel about 1 cm further into the CsI on average. The mean of the shower mean depth into the CsI was approximated by

$$z_e(E) = 0.11 + 0.018 \ln E, \quad (6.11)$$

$$z_\gamma(E) = 0.12 + 0.018 \ln E, \quad (6.12)$$

where E is the incident energy in GeV, and the result is in meters. The showers in the library recorded energy deposits in a 13×13 array of small crystals. The large crystals were treated as an 2×2 array of small crystals.

The entries in the EM shower library were segmented in energy in six bins; 2, 4, 8, 16, 32, and 64 GeV. Before lookup, the actual energy was smeared by a Gaussian shape to reproduce the observed energy resolution of the CsI. The energy bin for shower lookup included a finite probability for picking a neighboring bin. The probability was based upon a logarithmic interpolation between the neighboring bin energies and suppressed systematic effects due to bin boundaries.²

The entries were also binned according to position. Figure 47 shows the position bins defined on a small CsI crystal. The bins are 0.7 mm in the center of the crystal and 0.2 mm at the boundary since the dependence of energy sharing among crystals was more sensitive to position at the crystal boundaries. The position bins need only be defined for an octant of the crystal since the other shower configurations can be inferred from the symmetry. When an electron or photon is incident at the CsI, an electro-magnetic cluster is selected from a given position bin based on its XY position at the shower mean, given by equation 6.12.

The response was also divided into 25 2 cm. bins along the Z direction. The entries

²Specifically, the shower shape could change from bin to bin inducing small non-uniformities into the applied threshold correction [Prasad 98].




Figure 47. The position bins used to store GEANT generated electro-magnetic showers for the CsI simulation. Shower lookups for regions of the crystal outside of the indicated region can be related to the defined bins by symmetry.

from the library were convoluted with known Z dependencies in longitudinal light yield from each crystal as found by the cosmic ray muons, (see chapter 2.) Finally, the energy response is scaled to the energy of the incident particle from the library energy.

Electro-magnetic shower leakage out the back is not simulated, and was assumed to be absorbed in the Pb Wall.

6.1.9.2 Pion Showers

When a pion struck the front face of the CsI, an entry was fetched from the pion shower library. The pion shower library was compiled from GEANT simulations of pions striking the CsI and depositing energy in the CsI and downstream detectors. The shower library is organized into twelve momentum bins ranging from 4 GeV to 64 GeV. Like the EM shower simulation, the energy response is indexed by a position bin in one octant of a CsI crystal. However, there were only 10 divisions of X and Y across a crystal, each of 2.5 mm spacing. The x,y position of a pion at the CsI face determines which position bin

to look in when fetching a pion shower from the library. Each entry in the library also contained energy deposits for non-CsI detectors, including the HA.

6.1.9.3 *Muon MIPS*

When a muon hit the front face of the CsI, it was allowed to pass through and a MIP signal was generated in the intersecting crystals. The MIP signal was based on the Bethe-Bloch energy loss in dense materials with Landau fluctuations [Leo 87]. The muon was modeled to undergo multiple scattering in the CsI, and the signal was shared among crystals if the muon traversed more than one crystal. The muon, if still moving, was also scattered through the lead wall, and energy deposit to the HA was added.

6.1.9.4 *The HA*

The HA was filled with energy from the pion showers. Only the total sum of energy deposit from pion showers was kept during this stage. Eventually, the sum of deposited energy was compared to the HA threshold and, if surpassing, a trigger latch bit was set. Energy in individual HA counters was also recorded in the muon simulation.

6.1.9.5 *Conclusion*

Once energy deposit was completed in the CsI, the total energy E_T in the CsI was found and compared to thresholds. If E_t was above a given threshold, the appropriate trigger latch bit was set. The individual crystals were also compared to thresholds. The crystals above threshold were set as valid HCC bits. Also, the time structure of the PMT pulse in slices was simulated as well as the subsequent DPMT response based on known responses measured in the data.

6.1.10 Accidentals Simulation

Accidental events were recorded on tape using the accidental trigger as described in chapter 2. This trigger was assumed to be uncorrelated with detector activity due to kaon decays. After the complete generation and tracing of events through the detector, an accidental event is read from disk. The energies from this accidental event are overlaid on the simulated energies from the decay. In the drift chamber system, hits were overlaid on the generated hit list. If any generated hit followed an accidental hit, the generated hit was eliminated to simulated the discriminator deadtime.

It is important to stress again that the accidental energy was added before threshold evaluations and trigger evaluations were done; since accidentals could kill an event by firing a veto counter or cause an event to be accepted by raising a threshold energy.

6.2 Individual Decay Generators

6.2.1 Monte Carlo Simulation of $K_L \rightarrow 2\pi^0$

6.2.1.1 Strangeness Oscillations

In one of the most beautiful demonstrations of principle in all of physics, it can be shown that as a consequence of equations 1.7 and 1.8 and the differing decay lifetimes, the overall strangeness of a beam of neutral kaons will oscillate. Since the K_1 is not the antiparticle of the K_2 , their masses may differ while being consistent with the *CPT* theorem. This is furthermore a consequence of the fact that the K_1 and K_2 states are degenerate in the strong and EM interactions, but are split by the weak interaction. The measured value of the mass difference is 3.49×10^{-6} eV [Schwingerheuer 95].

The evolution of a kaon of strangeness $S = +1$ (pure K^0) at proper time $t = 0$ can then be described using the definitions of the K_1 and K_2 above.

$$|\Psi(t)\rangle = \frac{1}{\sqrt{2}} \left[|K_1\rangle e^{-im_1 t} e^{-\Gamma_1 t/2} + |K_2\rangle e^{-im_2 t} e^{-\Gamma_2 t/2} \right], \quad (6.13)$$

where $m_{1,2}$ is the mass and $\Gamma_{1,2}$ is the decay width of the $K_{1,2}$ respectively. This is

equivalent up to phase to

$$|\Psi(t)\rangle = \frac{1}{\sqrt{2}} \left[|K_1\rangle e^{-\Gamma_1 t/2} + |K_2\rangle e^{-i\Delta m t} e^{-\Gamma_2 t/2} \right]. \quad (6.14)$$

The probability of finding a K^0 in the beam at proper time t is given by

$$|\langle K^0 | \Psi(t) \rangle|^2 = \frac{1}{4} \left[e^{\Gamma_1 t} + e^{-\Gamma_2 t} + 2e^{-(\Gamma_1 + \Gamma_2)t} \cos \Delta m t \right]. \quad (6.15)$$

The presence of the interference term, in the absence of CP violation, cannot be seen in any of the kaon decay modes that are also CP eigenstates; although it can be inferred from interactions sensitive to the strangeness, such as the semi-leptonic decays of the kaon.³ However, it is a manifestation of CP violation that the interference term is observable in the $K \rightarrow 2\pi$ system. In the case of $2\pi^0$, of interest in the normalization analysis, the observed decay rate through the $\pi^0\pi^0$ decay channel for an initial K^0 particle as a function of proper time t is given by

$$|\langle \pi^0\pi^0 | H_w | \Psi(t) \rangle|^2 = \left[e^{\Gamma_S t} + |\eta_{00}| e^{-\Gamma_L t} + 2|\eta_{00}| e^{-(\Gamma_S + \Gamma_L)t/2} \cos(\Delta m t + \phi_{00}) \right], \quad (6.16)$$

where

$$\eta_{00} = \frac{A(K_L \rightarrow \pi^0\pi^0)}{A(K_S \rightarrow \pi^0\pi^0)}, \quad (6.17)$$

and ϕ_{00} is the phase of η_{00} .

The decay $K_L \rightarrow 2\pi^0$ is a CP violating decay. Thus, the observable $\pi^0\pi^0$ final state has contributions from both K_L and K_S . To account for this in the KTeV monte carlo, the full evolution of K_L and K_S amplitudes from the production target is taken into account. The Z distribution for each kaon is taken from expression 6.16. Since this is a 2-body decay, all of the decay product 4-vectors are constrained.

6.2.2 π^0 Dalitz Decay

The π^0 decay is dominated by the 2γ mode. Being an electro-magnetic decay, the distance traveled by a π^0 in our detector is extremely short. From [PDG 98], $c\tau$ for the

³This statement depends on the $\Delta S = \Delta Q$ rule, that the change in leptonic charge is equal to the change in strangeness, in all weak decays. The $\Delta S = \Delta Q$ rule has been tested by the CPLEAR collaboration [Angelopoulos 98] and is of fundamental importance in the interpretation of many experimental results.

π^0 is only about 25 nm. Therefore, π^0 decay is assumed to take place at the point of π^0 production.

The Dalitz decay of the π^0 is characterized by internal conversion of one of the photons in the more copious $\pi^0 \rightarrow \gamma\gamma$ decay and occurs with a branching ratio of about 1.2%. The rate has been calculated by Kroll and Wada [Kroll 55], and the KTeV monte carlo includes radiative corrections [Ter-Mikaelian 72] that takes into account the inner bremsstrahlung corrections like $\pi^0 \rightarrow e^+e^-\gamma\gamma$. The monte carlo also includes the form factor measured by the CELLO collaboration [Behrend 91].

6.2.3 Monte Carlo Simulation of $K_L \rightarrow \pi^0 e^+ e^- \gamma$

The decay generator for $K_L \rightarrow \pi^0 e^+ e^- \gamma$ in the KTeV monte carlo was adapted for use from a program supplied (very generously!) by F. Gabbiani. The original program performed a monte carlo integration of the $K_L \rightarrow \pi^0 e^+ e^- \gamma$ matrix element in order to calculate the branching ratio. This implementation of the $K_L \rightarrow \pi^0 e^+ e^- \gamma$ decay generator uses the same matrix element calculation to weight generated monte carlo events. The calculation is done for $K_L \rightarrow \pi^0 \gamma\gamma^*$ where the γ^* denotes the off-shell photon which subsequently converts internally to an e^+e^- pair. The variables used in the determination of the matrix element are

$$Z = \frac{M_{\gamma\gamma^*}^2}{M_K^2}, \quad (6.18)$$

$$Q = \frac{M_{\gamma^*}^2}{M_K^2}, \quad (6.19)$$

and

$$Y = \frac{E_\gamma - E_{\gamma^*}}{M_K}, \quad (6.20)$$

where M_K is the mass of the K_L , $M_{\gamma\gamma^*}$ is the invariant mass of the $\gamma\gamma^*$ system, M_{γ^*} is the mass of the e^+e^- pair, and E_γ and E_{γ^*} denote the energies of the on and off shell photons in the CM frame of the K_L . The matrix element takes the form

$$M.E.(Y, Z, Q) = \frac{\partial BR(K_L \rightarrow \pi^0 \gamma\gamma^*)}{\partial Z \partial Q \partial Y}(Y, Z, Q) \quad (6.21)$$

In order to generate a decay, we throw an event in random variables Z , Y , and Q and weight the event by the matrix element. In order to accept an event, a flat random




Figure 48. The decay distribution in Q as found by the KTeV Monte Carlo. This decay generator is based upon that used in [Gabbiani 97].

number r is thrown in the range $[0, \max(M.E.)]$ and if the weight exceeds r , the event is accepted.

To complete the event generation, we need to find the decay angle, θ , of the γ in the $\gamma\gamma^*$ CM frame. Using simple kinematical relations, we find that

$$\cos(\theta) = \frac{2YZ - (1 + Z - D)Q}{(Z + Q)\lambda^{\frac{1}{2}}(1, Z, D)}, \quad (6.22)$$

where

$$D = \frac{M_{\pi^0}^2}{M_K^2}, \quad (6.23)$$

and λ is the triangle function

$$\lambda(a, b, c) = a^2 + b^2 + c^2 - (ab + bc + ac). \quad (6.24)$$

Finally, we need to obtain a form for the e^+e^- pair from the γ^* . Since we already have the e^+e^- invariant mass, I use a modified Kroll-Wada form to generate the e^+e^-




Figure 49. The decay distribution in Y as found by the KTeV Monte Carlo. This decay generator is based upon that used in [Gabbiani 97].

decay angle. The full Kroll-Wada matrix element described the internal conversion of a photon in the Dalitz π^0 decay. Neglecting everything except the angular distribution, a weight is thrown according to the formula

$$W_{KW} = (1 + \cos(\theta)) + b_{KW}(1 - \cos(\theta)), \quad (6.25)$$

where

$$b_{KW} = \frac{4M_e^2}{M_{\gamma^*}}. \quad (6.26)$$

The Z position of the decay vertex for this decay mode is assumed to be from a K_L ; in other words, the Z vertex distribution is pure K_L exponential.




Figure 50. The decay distribution in Z as found by the KTeV Monte Carlo. This decay generator is based upon that used in [Gabbiani 97].

6.2.4 Monte Carlo Simulation of $K_L \rightarrow e^+e^-\gamma\gamma$

The decay generator for $K_L \rightarrow e^+e^-\gamma\gamma$ in the KTeV monte carlo was originally provided by Herb Greenlee and produced weighted monte carlo events. The calculation was based upon a direct calculation of the process $K_L \rightarrow e^+e^-\gamma\gamma$ and was used in the previous search for $K_L \rightarrow \pi^0 e^+e^-$ and the previous measurement of $K_L \rightarrow e^+e^-\gamma\gamma$ branching ratio from this collaboration.

Since the decay $K_L \rightarrow e^+e^-\gamma\gamma$ is so closely related to the decay $K_L \rightarrow e^+e^-\gamma$ a new monte carlo generator was devised which unites the treatment of the two decay modes. The events produced by this generator have been rolled against a random number so that they are unweighted.

6.2.5 Monte Carlo Simulation of $K_L \rightarrow 3\pi_D^0$

$K_L \rightarrow 3\pi_D^0$ was simulated with a flat phase space distribution. Since this mode is only considered as a background and since the effects of the measured form factor are small, this simulation was considered sufficient. Full amplitude evolution of the decay vertex into the decay region was done for this decay mode. However, it essentially reduces to the pure K_L exponential distribution since the rate of $K_S \rightarrow \pi^0\pi^0\pi^0$ is so small.

CHAPTER 7

The Normalization Mode Analysis

In order to normalize the measurement of the branching ratio of $K_L \rightarrow \pi^0 e^+ e^- \gamma$, we must also count the number of K_L particles that decayed in the decay volume during the run. This number is frequently referred to as the K_L *flux*. It is the flux measurement that comprises the *normalization mode* analysis. For the decays $K_L \rightarrow \pi^0 e^+ e^- \gamma$ and $K_L \rightarrow \pi^0 e^+ e^-$, we normalize to the decay $K_L \rightarrow 2\pi_D^0$. The branching ratio measurement is then

$$BR_S = \frac{N_S}{N_N} \frac{A_N}{A_S} \frac{1}{BR_N}, \quad (7.1)$$

where S, N denote the signal and normalization mode respectively, BR denotes the branching ratio, A denotes the acceptance measured in the monte carlo, and N denotes the number of events observed in the data. If N_S is a small number (as is the case for $K_L \rightarrow \pi^0 e^+ e^-$), then the method of Feldman and Cousins is used to find a 90% upper limit on the measurement [Feldman 98].

Other decay modes exist which could have been chosen for normalization in this analysis. For example, $K_L \rightarrow \pi^+ \pi^- \pi^0$ or $K_L \rightarrow e^+ e^- \gamma$ could have been used. However, $K_L \rightarrow e^+ e^- \gamma$ was not collected in the same trigger sample. Furthermore, $K_L \rightarrow \pi^+ \pi^- \pi^0$ is heavily biased in our trigger sample by cuts that bias against charged pions. Using either of these two modes for normalization then entails a detailed understanding inter-trigger biases or pion shower shapes in the CsI. While the decay mode $K_L \rightarrow 2\pi_D^0$ has a sizeable uncertainty in its branching ratio, it turns out that this is not an issue in the $K_L \rightarrow \pi^0 e^+ e^- \gamma$ branching ratio measurement because it is already statistics limited.




Figure 51. The $K_L \rightarrow 2\pi_D^0$ decay final state. Two photons come from a π^0 , while the remaining photon and e^+e^- pair also come from a π^0 which decays via the Dalitz decay mode.

Furthermore, $K_L \rightarrow 2\pi_D^0$ shares the same final state as $K_L \rightarrow \pi^0 e^+ e^- \gamma$ so that systematic effects arising from the different final states will be small.

The decay $K_L \rightarrow 2\pi_D^0$ is distinguished by a final state including an e^+e^- pair and three photons. Two of the photons come from a π^0 , and the remaining three particles also come from a π^0 which decayed through the Dalitz channel. (See figure 51.) This π^0 will be referred to as the Dalitz π^0 , or π_D^0 . The basic reconstruction requirements at this stage are

- 2ENCLUS trigger, 2ENCLUS L3 Tag, and Crunch Requirements (See Chapter 3)
- 2 charged tracks (See Chapter 5)
- 5 hardware clusters in the CsI (See Chapter 5)
- A decay vertex has been found (See Chapter 5)

Two of the five clusters must match to the electron tracks, and there are three pairing choices possible to assign the photons to the π^0 or to the Dalitz π^0 . The invariant mass of the π^0 is calculated for each of the three possible pairings assuming that they arise from a particle decay at the charged vertex. The *charged* vertex is determined using the trajectories of the two charged tracks, as determined by the upstream pair of drift chambers, and is the vertex returned by the tracking code. The pairing choice with the invariant mass closest to the nominal π^0 mass is chosen as the best pairing. The unpaired photon will be referred to as the lone photon. In $K_L \rightarrow 2\pi_D^0$ decay, the lone photon is

Kaon Mass Resolution




Figure 52. The $e^+e^-\gamma\gamma\gamma$ mass resolutions using charged and neutral vertex positions. Since the neutral vertex calculation implicitly constrains the $\gamma\gamma$ mass to be the π^0 mass, improved resolution results.

from a π^0 Dalitz decay; while in $K_L \rightarrow \pi^0 e^+ e^- \gamma$ decay, the lone photon is from direct emission.

After a pairing choice is made, the vertex Z position can be recalculated assuming that the two paired photons proceed from a π^0 . Using simple kinematics and small angle approximation,¹

$$\Delta Z = \Delta R_{12} \frac{\sqrt{E_1 E_2}}{m_{\pi^0}}, \quad (7.2)$$

where R_{12} is the distance between the two photons at the CsI face, $E_{1,2}$ are the photon energies, and ΔZ is the distance from the π^0 decay vertex to the CsI. The Dalitz π^0 mass and the overall eegg mass are calculated using this *neutral vertex*. Since the π^0 mass is constrained in calculations using the neutral vertex, improved mass resolutions result. Therefore, the neutral vertex masses are used for mass cuts.

In the following sections, the cuts to isolate a clean $K_L \rightarrow 2\pi_D^0$ sample in E799 will be presented.

¹The maximum angle involved is about 0.14 radians.

	π^0 Mean	π^0 Resolution	π_D^0 Mean	π_D^0 Resolution
Winter Data	135.62 ± 0.01	2.42 ± 0.03	134.76 ± 0.01	1.88 ± 0.02
Winter MC	135.42 ± 0.01	2.43 ± 0.02	134.84 ± 0.01	1.87 ± 0.01
Summer Data	135.60 ± 0.01	2.40 ± 0.03	134.72 ± 0.01	1.90 ± 0.02
Summer MC	135.42 ± 0.01	2.42 ± 0.02	134.84 ± 0.01	1.84 ± 0.01
Winter Corrected	135.56 ± 0.01	2.38 ± 0.02	134.82 ± 0.01	1.88 ± 0.02
Summer Corrected	135.61 ± 0.01	2.38 ± 0.02	134.81 ± 0.02	1.86 ± 0.02

Table 15. Mass resolutions for the various π^0 masses calculated in the normalization mode analysis. The mass peaks in each case were fit to a single Gaussian in a close range about each peak. In all cases, the fit χ^2 was between 1.5 and 2.0, and the quoted errors are the fit errors. The quoted resolution is equal to the sigma of the fit Gaussian.

7.1 Cuts To Isolate $K_L \rightarrow 2\pi^0, \pi^0 \rightarrow e^+e^-\gamma$

In order to identify $K_L \rightarrow 2\pi_D^0$ events, we place cuts on the invariant masses of the $\gamma\gamma$ $e^+e^-\gamma$ and $e^+e^-\gamma\gamma\gamma$ in the best pairing. The mass resolutions of the various mass distributions were evaluated after all other cuts were made, and appear in table 15 and 16. While the mean masses disagree with the PDG² values for the masses by as much as 0.3%, this difference is reflected in both the data and the monte carlo so that the systematic error associated with cuts on these quantities should be small. The “corrected” values in the tables reflect data analyzed using the special run by run CsI corrections and the CsI non-linearity fudge, described in chapter 5. While these corrections improve the π^0 resolution slightly, the mean differences from PDG masses remain. In any case, the resolutions match very well. Figure 53 shows the π^0 and π_D^0 mass peaks in the data after all cuts with monte carlo overlay. The mass cuts are

- $127 < M_{\gamma\gamma} < 143 \text{ MeV}/c^2$
- $129 < M_{e^+e^-\gamma} < 141 \text{ MeV}/c^2$

²PDG : *Particle Data Group*, the maintainer of the official best values for particle physics measurements.

	$e^+e^-\gamma\gamma\gamma$ Mean	$e^+e^-\gamma\gamma\gamma$ Resolution
Winter Data	497.57 ± 0.02	3.41 ± 0.02
Winter MC	497.42 ± 0.02	3.37 ± 0.02
Summer Data	497.76 ± 0.02	3.44 ± 0.03
Summer MC	497.39 ± 0.02	3.38 ± 0.02
Winter Corrected	497.31 ± 0.02	3.41 ± 0.02
Summer Corrected	497.28 ± 0.03	3.46 ± 0.03

Table 16. Mass resolutions for $e^+e^-\gamma\gamma\gamma$ mass calculated in the normalization mode analysis. The mass peaks in each case were fit to a single Gaussian in a close range about each peak. In all cases, the fit χ^2 was between 1.5 and 2.0, and the quoted errors are the fit errors. The quoted resolution is equal to the sigma of the fit Gaussian.




Figure 53. The π^0 and π_D^0 mass peaks in the winter E799 dataset for the $K_L \rightarrow 2\pi_D^0$ sample after all cuts with monte carlo overlay. The resolutions match very well, and the means between data and monte carlo agree to within 0.15%.

- $487 < M_{e^+e^-\gamma\gamma\gamma} < 507 \text{ MeV}/c^2$

and correspond roughly to 3 sigma cuts about the mean.

Finally, a cut is made on the decay volume. The decay vertex³ is required to be between 95 m and 158 m. The low side cut removes decays that are disturbed by the fringe fields of the final sweeper magnet, and the high side cut removes decays that originate as conversions in the vacuum window.

- $95 < V_Z < 158 \text{ meters from target}$

At this point, we might expect to have a fairly clean and background free sample of $K_L \rightarrow 2\pi_D^0$ events. However, special backgrounds exist which are not rejected by any of the cuts discussed so far. The background is from $K_S \rightarrow \pi^0\pi_D^0$. There are two ways in which K_S decays contribute to the background in $K_L \rightarrow 2\pi_D^0$, both of which are rather amusing. The other major threat to the purity of the normalization sample is from $K_L \rightarrow 3\pi_D^0$.

7.1.1 K_S Backgrounds

7.1.1.1 *Collimator Scatters*

The first amusing source of K_S contamination comes from K_L particles which scatter in the inner edges of the defining collimator. A background is obvious when we look at the transverse momentum of events in the decay sample. The square of the transverse momentum, P_T^2 , is a measure of the quality of reconstruction of an event. Simply put, if an event is well reconstructed, the the total reconstructed momentum vector had better be parallel to the line connecting the center of target with the decay vertex. P_T is a measure of the reconstructed momentum not parallel to this vector.

$$P_T^2 = |\vec{P} - (\hat{V}\vec{P})\hat{V}|^2. \quad (7.3)$$

³It is helpful to keep in mind that for the π^0 Dalitz modes, the opening angle between the two electrons tends to be small. Thus, in the $K_L \rightarrow 2\pi_D^0$ mode, the neutral vertex has better vertex position resolution than the charged vertex position. In the signal modes, the charged vertex is better. Therefore, when cutting on Z vertex position, cuts are placed on both charged and neutral vertex positions.




Figure 54. Collimator Scatters. The top plot shows the log distribution of transverse momentum for E799 $K_L \rightarrow 2\pi^0$ data after all cuts. The filled circle overlay is the $K_L \rightarrow 2\pi^0$ monte carlo prediction with no collimator scatters, and the filled triangle overlay is the collimator scatter ONLY monte carlo prediction. Kaons scattering in the collimator tend to regenerate to K_S so that the fraction of collimator scatters into 2π is relatively high. The cut shown is the final analysis cut in P_T to remove the collimator scatters. The bottom plot shows the extrapolated XY position of the kaon in the collimator plane for events which fail the P_T cut. The resulting plot is aptly called the “hotlips” plot.

Figure 54 (top) shows a plot of the $|P_T|$ distribution in the normalization sample plotted on a log scale, data with monte carlo overlay. The peak at small $|P_T|$ is well predicted by the normalization mode monte carlo. However, a huge peak also exists at large P_T^2 that is not predicted by the regular monte carlo. Studying these events exclusively, it was found that these events arise from scattering in the collimator. K_L events that scatter in

the collimator material tend to regenerate to K_S , which subsequently decay to $2\pi^0$. The decay Z distribution of these events is exponential falling in our decay volume, supportive of the K_S hypothesis. The conclusive evidence comes from the XY position of the kaons at the Z position of the defining collimator. The “hot lips” are clearly visible in 54.

A simple cut on the P_T^2 distribution is enough to remove this background. The remaining background due to K_L to K_S collimator scatters is estimated to be less than 0.15%.

7.1.1.2 Target K_S

Another way in which K_S can contribute to the background in the normalization mode is by simply surviving from the target. Both K_L and K_S are produced in the beam, and $K_S \rightarrow 2\pi$ is a major K_S decay mode. Some K_S survive in the decay sample, and these tend to populate the sample at high momentum and at the upstream end of the decay volume. Since the K_L decay is CP suppressed, the amount of K_S in the final $2\pi^0$ sample is as high as a few percent. This *target K_S* background can in principle be removed by making a cut on the ratio of vertex Z position to momentum, Z/P , which is turn proportional to the number of lifetimes from target. However, the target K_S 's are well understood and are predicted by the monte carlo using the kaon oscillation dynamics presented in chapter 6, (see equation 6.16.) Since the monte carlo fully models the K_L - K_S dynamical evolution, the individual K_L and K_S amplitudes in the monte carlo can be used to estimate the fraction of K_L in the final data sample. To get the monte carlo K_L fraction, one weights the monte carlo events by the ratio of K_L amplitude at the decay point to total amplitude,

$$W_{K_L} = \frac{|\eta_{00}\alpha_L|^2}{|\alpha_S + \eta_{00}\alpha_L|^2}, \quad (7.4)$$

where W is the weighting factor, $\alpha_{L,S}$ is the $K_{L,S}$ amplitude at the decay point following evolution, and a η_{00} is the ratio of K_L to K_S amplitude into $2\pi^0$ (equation 6.17.) To obtain a K_L flux, one merely uses the sum of generated W_{K_L} in place of the sum of generated events when calculating the normalization mode acceptance, (see equation 6.2.) Also, one can estimate the amount of K_S decays in the final sample using the same method,

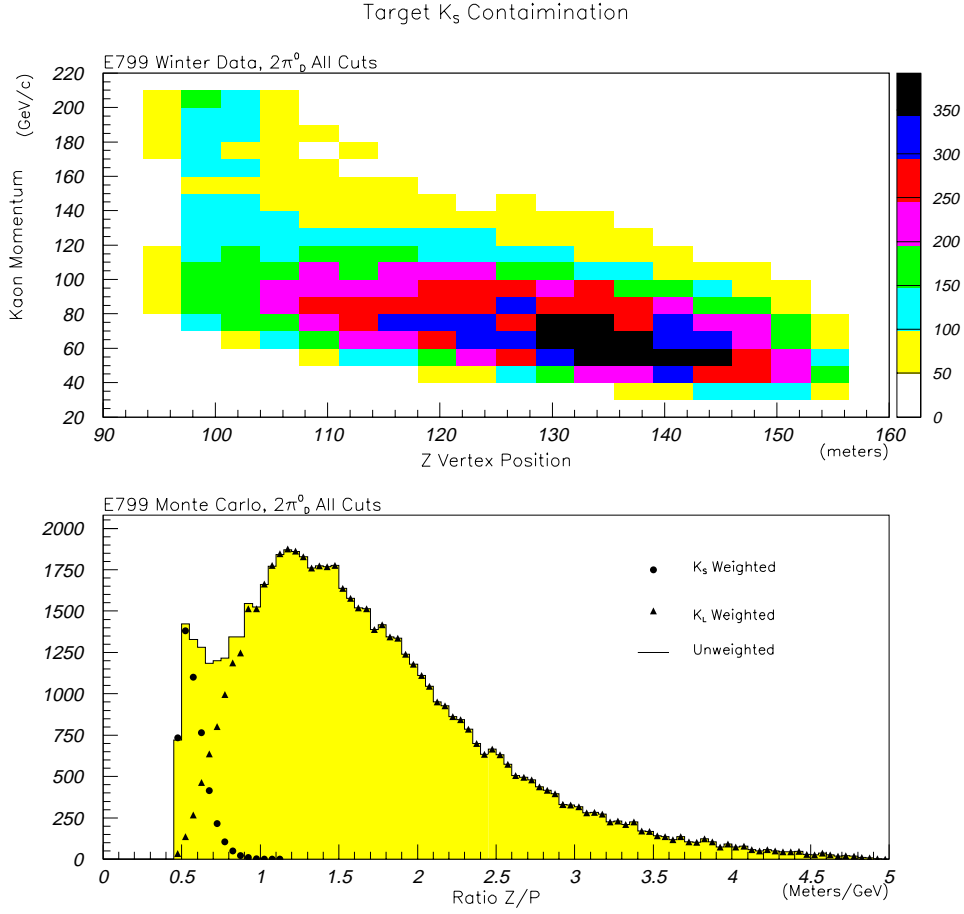


Figure 55. The top plot shows the distribution of the final E799 $K_L \rightarrow 2\pi_D^0$ data in vertex Z position and kaon momentum. The region at high momentum closer to the target is populated by decays of K_S that survive from the target. The bottom plot shows the Z/P ratio for $K_L \rightarrow 2\pi_D^0$ monte carlo events. The K_S survivors from the target dominate the region $Z/P < 0.75$. The filled circle overlay are the same monte carlo events weighted to be K_S , and the filled triangle overlay are the same events weighted to be K_L according to equation 7.4.

obtaining about 5%. For the signal decay mode, the decay distribution is assumed to be pure K_L so that no weighting is necessary.

Figure 55 (top) shows the distribution of final $K_L \rightarrow 2\pi_D^0$ events plotted as a scatter-plot in momentum versus vertex Z position. Figure 55 (bottom) shows the distribution of $K_L \rightarrow 2\pi_D^0$ monte carlo events in Z/P with overlays using K_L weighted events and K_S weighted events.




Figure 56. $e^+e^-\gamma\gamma\gamma$ invariant mass is plotted versus \log of $|P_T|$. $K_L \rightarrow 3\pi_D^0$ background as well as collimator scatters are well excluded by the indicated cuts in mass and P_T . Also, a popular corporate logo.

7.1.2 $K_L \rightarrow 3\pi_D^0$ Background

Figure 56 shows a plot of $e^+e^-\gamma\gamma\gamma$ invariant mass versus \log of $|P_T|$ for the normalization mode sample. This artful plot clearly illustrates all of the major backgrounds in the normalization mode. mode.

Fortunately, no special cuts need to be invented against the $K_L \rightarrow 3\pi_D^0$ sample at this stage. Using the monte carlo, the $K_L \rightarrow 3\pi_D^0$ background represents less than 0.05% percent of the normalization sample at this point, estimating from the sidebands. We can safely ignore the $K_L \rightarrow 3\pi_D^0$ background here.




Figure 57. The $e^+e^-\gamma\gamma\gamma$ final state mass for the $K_L \rightarrow 2\pi_D^0$ normalization mode after all cuts for the winter E799 dataset with monte carlo overlay. The background at low invariant mass below the kaon peak is from $K_L \rightarrow 3\pi_D^0$ fragments. In the mass window between 487 and 507 MeV/c^2 , there are estimated to be about 10 $K_L \rightarrow 3\pi_D^0$ events based upon a fit of the background shape in the low mass sideband.

7.2 Cuts To Improve Data Quality

Several cuts have been made to exclude edges in the detector, to remove possible threshold effects, and to verify trigger signals. These are listed below.

- A cut was made to require CA energy below 12 GeV to verify the trigger threshold.
- A cut was made 1.25 cm around the CsI beam holes to reduce systematic effects from a small data/MC disagreement around the beam holes.
- E/p is required to be between 0.95 and 1.05. This is an electron quality cut.
- A simulation the Etotall trigger energy, ETsim, is constructed using raw block energies in the CsI. ETsim is required to be above 30 GeV to verify the Etotall trigger threshold. Also, a small data/MC disagreement is removed.
- The VV' trigger counters are checked to make sure that the trigger should have fired.
- The minimum cluster energy is required to be above 2.0 GeV. This is done to remove systematic effects from a small data/MC disagreement and assures us that the HCC bit threshold is satisfied.
- The kaon momentum is required to be less than 210 GeV/c.
- If the ee invariant mass, m_{ee} , is greater than 0.14 GeV/ c^2 then the event is rejected. For the $K_L \rightarrow 2\pi_D^0$ decay, it is kinematically impossible to have $m_{ee} > 0.14$ GeV/ c^2 .
- The event is rejected if the separation between the two tracks is less than 1 cm at DC1 in X and separately in Y. This cut is made to eliminate small data/MC disagreements below 1 cm.

7.3 Flux Results and Systematic Error

With the set of cuts outlined above, the flux results are shown in table 18. Before

Run Period	$K_L \rightarrow 2\pi_D^0$ Acceptance	Number Observed
Winter (before 8245)	$(0.898 \pm 0.008)\%$	2256
Winter (after 8245)	$(0.947 \pm 0.004)\%$	30536
Summer	$(0.971 \pm 0.004)\%$	23675

Table 17. The acceptances and numbers of observed $K_L \rightarrow 2\pi_D^0$ events in E799. The winter run is split at run 8245 to take account of a slight acceptance loss before run 8245 due to a swapped HCC cable.

exploring detector based systematics, there is a systematic error associated with the existing branching ratio measurement of $K_L \rightarrow 2\pi_D^0$. The $K_L \rightarrow 2\pi^0$ branching ratio is [PDG 98]

$$BR(K_L \rightarrow 2\pi^0) = (9.36 \pm 0.2) \times 10^{-4}, \quad (7.5)$$

with an error of 2.13%. The π^0 Dalitz branching ratio is

$$BR(\pi^0 \rightarrow e^+e^-\gamma) = (1.198 \pm 0.032)\%, \quad (7.6)$$

with an error of 2.67%. The combined branching ratio is

$$BR(K_L \rightarrow \pi^0\pi^0; \pi^0 \rightarrow e^+e^-\gamma) = (2.24 \pm 0.08) \times 10^{-5}, \quad (7.7)$$

with an error of 3.42%, where either π^0 can decay via the Dalitz channel. This error will be counted as a systematic in the flux calculation, and eventually in the calculation of $BR(K_L \rightarrow \pi^0 e^+ e^- \gamma)$.

7.3.1 Data - Monte Carlo Comparisons

Using the $K_L \rightarrow 2\pi_D^0$ sample, various systematic studies of the detector can be performed. Figures 58 and 59 show the track illuminations in DC1 in the winter and summer datasets with monte carlo overlay. As a measure of the illumination of the neutral beams, the ratio of vertex X and Y position to the Z position is plotted in figure 60 with monte

Run Period	Flux ($\times 10^{11}$)
Winter (before 8245)	0.112 ± 0.002
Winter (after 8245)	1.440 ± 0.008
Summer	1.088 ± 0.007
Total	2.64 ± 0.01

Table 18. The E799 K_L flux. The winter run is split at run 8245 to take account of a slight acceptance loss before run 8245 due to a swapped HCC cable. The quoted errors are statistical only.




Figure 58. X position illumination of tracks at DC1 in the summer and the winter with monte carlo overlay.




Figure 59. Y position illumination of tracks at DC1 in the summer and the winter with monte carlo overlay.

carlo overlay. These distributions match data to monte carlo fairly well, and the systematic effects associated with the small differences around the edges of these plots are small.

The kaon momentum distribution with monte carlo overlay is shown in figure 62. There is a slight slope associated with the data/mc ratio plot in the momentum distributions. The momentum cut was tightened to the range 40 to 160 GeV/c, and this leads to a flux change of +0.4% in the winter and +0.6% in the summer.

Finally, the distribution of hits in DC1,2 beam regions with monte carlo overlay is shown in figure 63. The number of beam hits is defined as the number of tracks intersections per event with a beam region in DC1 or DC2. The purpose is to try and

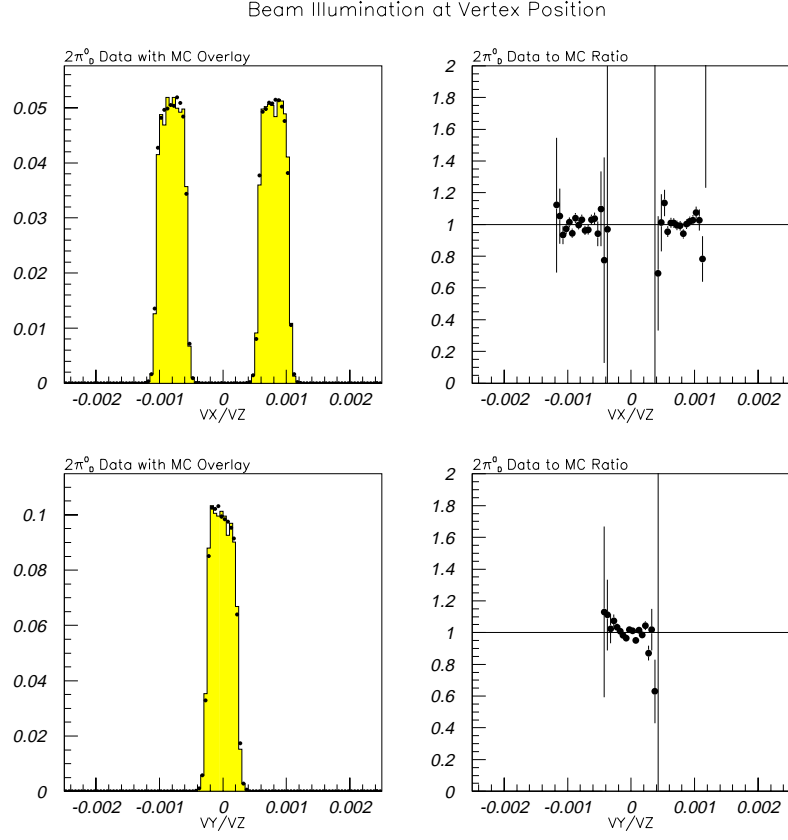


Figure 60. As a measure of the beam illumination, the ratio of vertex X and Y positions to the vertex Z position are plotted for both the summer and winter datasets with monte carlo overlay. The slope in the V_Y/V_Z distribution is due to the targeting angle.

find any inefficiencies in the beam regions of DC1 and DC2 due to high rate or radiation damage. No inefficiency is found in the winter E799 data; however, a small constant inefficiency is found in the summer data. This leads to a systematic error of +0.6% in the summer.

To estimate systematic error from the mass cuts, I used two methods. First, the mass cuts were moved in the data by an amount equal to the difference between the mass peaks in data and monte carlo. Each case, π^0 mass cut, π_D^0 mass cut, and kaon mass cut were evaluated separately and then altogether. The maximum deviation is taken to be the systematic error estimate. This occurred in the π^0 mass cut, and the change was +0.2%




Figure 61. Cluster X and Y position illumination at the CsI for both the summer and winter datasets with monte carlo overlay.

in the winter and in the summer; however, there were also variations that produced a decrease in flux by almost as much. Secondly, I re-did the analysis using the special CsI calibrations to correct for run by run energy scale variations and the non-linearity fudge, described in chapter 5. The change in the flux was +0.6% in the winter and in the summer.

Other systematic errors include

- The CA energy cut was tightened from 14 GeV to 1 GeV. This has the effect of chopping off the tail of the CA energy distribution. This changes the flux measurement by -1.3% overall and could be due to an excess of accidental energy near the CsI beamhole.




Figure 62. Kaon momentum distribution in the winter and summer.

- The cut on track separation at DC1 was dropped. This changes the flux measurement by -0.3% overall.
- The P_T^2 cut was changed from 1000 (MeV/c)^2 to 500 (MeV/c)^2 . This changes the flux measurement by +0.5%.
- The XT calibrations were fudged so that the SOD means matched between data and monte carlo. This is discussed in the next section. Including the XT fudge changes the flux measurement by +1.5%.

Including the systematic discussed below in section 7.3.2, the XT map and DC resolution fudge, these are the most important detector based systematics. In addition, each cut was varied and the systematic is recorded in table 19.

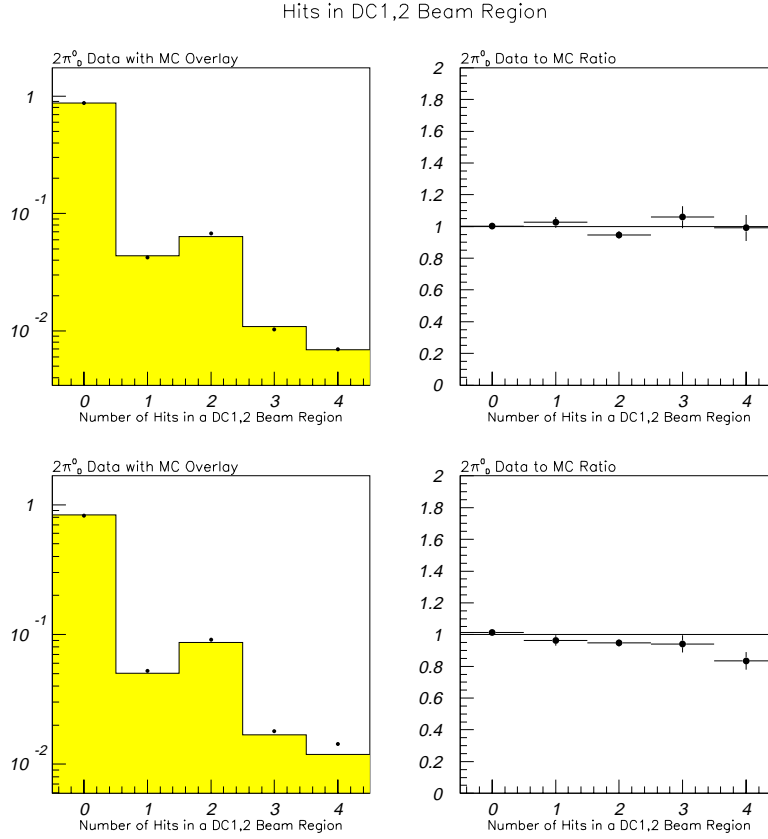


Figure 63. The distribution of number of beam hits per event in DC1 and DC2 with monte carlo overlay. The top pair of plots show the overlay for E799 winter data, and no evidence for beam region inefficiency is found. The bottom pair of plots show the overlay for the E799 summer data, and a small constant inefficiency per beam hit is seen.

7.3.2 Plots Over Which I Lose Sleep

At this point, there is one remaining detector based systematic error which potentially dwarfs all the others. The data/MC comparison of the charged vertex Z position shows a systematic mismeasurement of the decay vertex position that is significant at the 2.5 to 3 sigma level.⁴ This is manifest in the data/MC ratio plots, (see the top plots of figures 71

⁴A similar slope exists in the neutral vertex position. Since the charged vertex depends only on tracking parameters, it is used for this systematic study rather than the neutral vertex position which depends upon calibrations from multiple detector systems.

Systematic	Total Flux Change
Tighten CA cut to 1 GeV	-1.3%
Corrected CsI Energy	+0.6%
Tighter Momentum Cut	+0.5%
Tighten P_T^2 cut	+0.5%
Fudge XT Maps and DC Res	-0.4%
Drop DC1 Track Separation Cut	-0.3%
Beam Hits in DC1,2	+0.25%
Varied Mass Cuts	$\pm 0.2\%$
Loosen E/p cut to 0.9-1.1	+0.2%
Total	1.7%

Table 19. A table of systematic errors in the normalization mode analysis in order of importance. The total estimated systematic error in the flux measurement is 1.7%.

and 7.3.2.) In the E799 $K_L \rightarrow 2\pi_D^0$ sample, the ratio of data to monte carlo of the vertex Z position has a slope of about $(0.12 \pm 0.04)\%$ per meter in the winter and $(0.11 \pm 0.04)\%$ per meter in the summer.

Many hypotheses have been advanced to explain the “Z slope” problem. Suspiciously, the Z slope problem tends to be associated with events where the opening angle between the two tracks is small. Figure 64 shows the dependence of the Z slope on opening angle in the copious $K_L \rightarrow 3\pi_D^0$ decay mode. Furthermore, a comparison of the Z slope problem among different analyses has shown that this is not a problem in modes that typically have higher opening angles between tracks or more than two tracks, such as $\pi^+\pi^-\pi^0$, or $e^+e^-e^+e^-$. This makes it a problem for the normalization $K_L \rightarrow 2\pi_D^0$ mode. In contrast, it is not likely to be as severe a problem in $K_L \rightarrow \pi^0 e^+ e^- \gamma$. Decays with smaller opening angles are more affected by this systematic effect because the uncertainty in the vertex position is higher in these cases.

This is an important problem. When cuts are adjusted at the edges of the decay volume, the flux result only changes by a few tenths of a percent. However, the flux

99/02/03 07.26

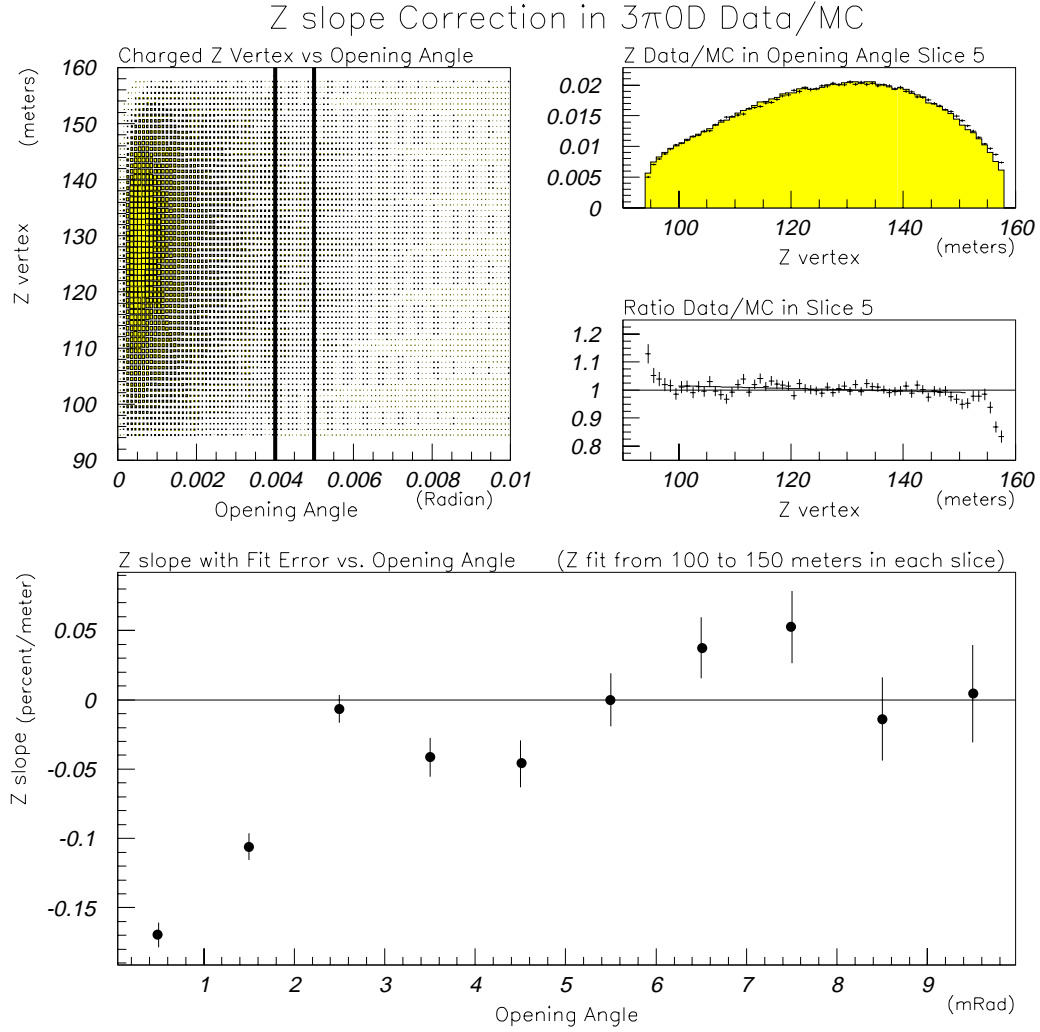


Figure 64. The Z slope problem in $K_L \rightarrow 3\pi^0$. In order to plot the dependence of the Z slope problem on opening angle, a 2 dimensional scatterplot (upper left) is filled with data from the plentiful $K_L \rightarrow 3\pi^0$ decay mode. The Z vertex position is plotted versus opening angle between the two electron tracks. Note that for this mode, like all π^0 modes, the plot is dominated by low opening angles. The scatterplot is sliced in opening angle, and the vertex Z position distribution is compared between data and MC for each slice (upper right). The ratio plot in each slice is fit to a line. Finally, the slope is plotted versus the opening angle (bottom plot.) ($K_L \rightarrow 3\pi^0$ data provided courtesy of Pat Toale.)




Figure 65. The SOD (sum of distances) distribution for plane pair DC2X in the winter E799 data from the $K_L \rightarrow 2\pi_D^0$ sample. The right side shows the data with monte carlo overlay, and the left side shows the ratio plot. The ratio plot shows the mismatch of the widths as a parabolic curve. The means of the two distributions is also offset from one another.

calculated using the upstream half of the decay volume (from 195 to 120 meters) is 2.77×10^{11} and that using the downstream half is 2.51×10^{11} kaons. The difference is $\pm 4.5\%$ overall. This number is regarded as an upper limit to the real systematic error, and it remains a priority to show that the flux estimation is better than that. It is also known that by cutting hard on vertex Z position error, one can also make the Z slope go away. However, this naturally biases against low opening angle events and events with low e^+e^- invariant mass. Such studies do show independently that the flux estimate given above changes by only -1.1% after making such a cut with an acceptance loss of 30% in the normalization mode. However, I was not willing to live with the loss of normalization mode acceptance nor with the biases induced by this cut in the signal modes.

Though it is likely to be the manifestation of several problems, I have endeavored in this analysis to explain the Z slope in terms of the basic tracking calibration constants in the experiment. Remember from chapters 2 and 5 the role of the XT maps in the track finding process. A critical check of the XT calibration is that the mean of the SOD distributions should be equal to 1/2 of a drift cell. Incredibly, this important check does not have input into the XT calibration itself. A data/MC comparison of typical

	X (m)	Y (m)
DC1 Winter	$\pm 0.090, \pm 0.167$	± 0.038
DC2 Winter	$\pm 0.082, \pm 0.173$	± 0.040
DC1 Summer	$\pm 0.082, \pm 0.176$	± 0.047
DC2 Summer	$\pm 0.085, \pm 0.182$	± 0.049

Table 20. The X and Y positions of the DC regions intersecting the beam in the winter and the summer for drift chambers 1 and 2.

SOD distributions appears in figure 65. The figure shows that not only do the SOD means disagree, but that the SOD width is also modeled poorly⁵. The SOD mean and width mismatches were cataloged by drift chamber and region within each drift chamber separately in the winter and the summer. A drift chamber region was defined as follows, (see figure 66). If the plane pair is in the X view, then there are five regions defined by the two strips containing beam holes and the three surrounding regions. In the Y view, there is only one strip containing both beam holes, and two surrounding regions. The X view regions are numbered from 1 to 5, and the Y view from 6 to 8. Each chamber is divided in this way to isolate the beam regions, which suffer from much higher rates. The beam regions used are defined in table 20. Figures 67 through 70 show the SOD mean and width mismatch in all chamber planes for the winter and the summer datasets in each of the above defined regions. The SOD means are typically lower in the data by 10 - 15 microns (about 0.2%), and the resolutions are typically off by 5%. The most important of the two effects turns out to be the deviation of the SOD means. A shift of SOD means by 0.2% can improve the Z slope problem dramatically; while an across the board shift of 10% in the resolutions is required to have the same effect on the Z slopes, but there is no physical basis for such a resolution shift.

The monte carlo SOD means were correct, at 6.35 mm across the board. In order to estimate systematic error from these effects, the XT maps were fudged by the amount

⁵The very high SOD tail, as discussed in chapter 5, is also modeled poorly. However, it turns out that modeling the SOD means and widths correctly is sufficient for our purposes.




Figure 66. The cartoon picture above shows the definition of drift chamber regions for a drift chamber. If the plane pair is in the X view, then there are five regions defined by the two strips containing beam holes and the three surrounding regions. In the Y view, there is only one strip containing both beam holes, and two surrounding regions. The X view regions are numbered from 1 to 5, and the Y view from 6 to 8. Each chamber is defined in this way to isolate the beam regions, which suffer from much higher rates.

shown in the plots, depending on the region traversed by each track, to make the data and monte carlo SOD means agree. A correction was also applied to the intrinsic chamber resolutions to correct these effects in each plane and region. The resulting Z slopes with the XT and resolution fudges are shown in figures 71 and 7.3.2. Using the corrected monte carlo, it is found that the flux measurement changes by -0.4%. The new Z slope is essentially gone in the summer, and is reduced to $(0.074 \pm 0.05)\%$ per meter in the winter.

While it is realized that any true corrections to the calibration constants must take run by run variations into account, the corrections applied above are valid within the limit of the statistics in this mode. A re-calibration of the drift chamber system in E799 is underway at the time of this writing. It is expected that the Z slope will disappear with better calibrations.

An alternate model assumes that the Z slope is due only to beam region inefficiencies in the drift chambers [KTeV 608]. This model also able to improve the Z slope in conjunction with a cut on the opening angle. The systematic error to be associated with the Z slope




Figure 67. The SOD means and widths in each DC region were fit to a Gaussian in a region close to the SOD peak in data and monte carlo. The above plot shows the ratio of fit SOD widths in the data to those in the monte carlo for the winter E799 run in $K_L \rightarrow 2\pi_D^0$. For a significant part of the winter run, the SOD widths were much higher in chamber 2 which, in addition to the other deviations, suggests a problem with the DC plane resolution modeling.




Figure 68. The SOD means and widths in each DC region were fit to a Gaussian in a region close to the SOD peak in data and monte carlo. The above plot shows the ratio of fit SOD means in the data to those in the monte carlo for the winter E799 run in $K_L \rightarrow 2\pi_D^0$. The SOD means were found to be lower by about 10 - 15 microns in the data. This suggests a systematic miscalibration of the XT maps.




Figure 69. The SOD means and widths in each DC region were fit to a Gaussian in a region close to the SOD peak in data and monte carlo. The above plot shows the ratio of fit SOD widths in the data to those in the monte carlo for the summer E799 run in $K_L \rightarrow 2\pi_D^0$. The deviations suggest a problem with the DC plane resolution modeling.




Figure 70. The SOD means and widths in each DC region were fit to a Gaussian in a region close to the SOD peak in data and monte carlo. The above plot shows the ratio of fit SOD means in the data to those in the monte carlo for the summer E799 run in $K_L \rightarrow 2\pi^0$. The SOD means were found to be lower by about 10 - 15 microns in the data. This suggests a systematic miscalibration of the XT maps.




Figure 71. The Z slope problem in the E799 winter $K_L \rightarrow 2\pi_D^0$ sample. The top plots show the vertex Z distribution with monte carlo overlay and the corresponding data/MC ratio with no special correction. The bottom plots show the same after a correction has been applied to the XT maps and the chamber resolutions.

problem estimated by this method is 1.8%, similar to the error found above. In any case, it is smaller than the systematic error from $\text{BR}(K_L \rightarrow 2\pi_D^0)$ alone.

7.4 Conclusion

The E799 K_L flux from $K_L \rightarrow 2\pi_D^0$ analysis is

$$\text{Flux} = (2.64 \pm 0.01_{\text{STAT}} \pm 0.04_{\text{SYS}} \pm 0.09_{\text{BR}}) \times 10^{11}, \quad (7.8)$$




Figure 72. The Z slope problem in the E799 summer $K_L \rightarrow 2\pi_D^0$ sample. The top plots show the vertex Z distribution with monte carlo overlay and the corresponding data/MC ratio with no special correction. The bottom plots show the same after a correction has been applied to the XT maps and the chamber resolutions.

assuming bad spill definitions for the $K_L \rightarrow \pi^0 e^+ e^- \gamma$ analysis and where the systematic error from the branching ratio has been listed separately.

CHAPTER 8

Analysis of $K_L \rightarrow \pi^0 e^+ e^- \gamma$

The final state of $K_L \rightarrow \pi^0 e^+ e^- \gamma$ includes an $e^+ e^-$ pair and three photons, two of which are distinguished by coming from a π^0 . The last photon comes directly from the K_L decay. The final state is shown in figure 73. The identification strategy begins when we identify the decay vertex using the charged tracks. The invariant masses of the three possible pairings of photons are calculated assuming that they are the result of a particle decaying to two photons at the charged decay vertex. The pairing with invariant mass closest to the nominal π^0 mass is assumed to come from the pion, and the remaining photon in the decay is assumed to be the lone photon. Lastly, the final state mass must match the nominal kaon mass. The invariant mass of the $e^+ e^- \gamma$ system is of theoretical interest, and we will therefore attempt to preserve as much of the spectrum as possible during the measurement while eliminating background from $K_L \rightarrow 2\pi_D^0$. The decay $K_L \rightarrow \pi^0 e^+ e^- \gamma$ is distinguished from $K_L \rightarrow \pi^0 \gamma \gamma$ only by the internal conversion of a direct photon of the latter decay to an $e^+ e^-$ pair. We therefore expect that the invariant mass spectrum of the $e^+ e^- \gamma$ will be similar to the invariant mass spectrum of the two direct photons in $K_L \rightarrow \pi^0 \gamma \gamma$ decay and the decay spectrum calculated in [Donoghue 97] and [Gabbiani 97].

Two copious backgrounds exist to this decay mode as well as several more subtle backgrounds. The decay $K_L \rightarrow 2\pi_D^0$ has the same particles in the final state as $K_L \rightarrow \pi^0 e^+ e^- \gamma$ and is distinguished by the fact that the $e^+ e^- \gamma$ part of the final state also comes from a π^0 . (See figure 51.) To eliminate this background, we make a cut against




Figure 73. The $K_L \rightarrow \pi^0 e^+ e^- \gamma$ decay final state. In $K_L \rightarrow \pi^0 e^+ e^- \gamma$ decay, two photons are the products of a π^0 decay while the third photon is from direct emission. The $e^+ e^-$ pair comes from an off-shell photon which converts internally. The decay dynamics are similar to those governing $K_L \rightarrow \pi^0 \gamma \gamma$ decay.

the π^0 mass in the $e^+ e^- \gamma$ system. The decay $K_L \rightarrow 3\pi_D^0$ can fake a $K_L \rightarrow \pi^0 e^+ e^- \gamma$ decay if two of the photons are missing or fused with other photons in the calorimeter. A cut on cluster shape in the CsI calorimeter effectively removes the danger from fusions. Cuts on mass and P_t can remove much of the remaining background, although their efficacy depends on which two photons are lost. Also, accidental photons can cause the fragmentary final state to appear to satisfy various mass cuts. Therefore, further cuts are made on kinematical variables to remove this background as well as additional cuts in the photon vetoes. See figure 74 for an example of how $K_L \rightarrow 3\pi_D^0$ can contribute to $K_L \rightarrow \pi^0 e^+ e^- \gamma$ background.

Charged pion backgrounds from decays such as $K \rightarrow \pi^\pm e^\mp \nu$ and $K_L \rightarrow \pi^+ \pi^- \pi^0$ with accidentals are small because of the electron identification cut on E/p in the CsI calorimeter. Although not used explicitly in this analysis, the TRD system has been used to check this assumption of small charged pion backgrounds.

Other backgrounds exist which come from photon conversions in the vacuum window or external bremsstrahlung from the vacuum window. Conversion and bremsstrahlung backgrounds originating from known decays, such as $K_L \rightarrow 2\pi_D^0$ or $K_L \rightarrow 3\pi_D^0$, are modeled by the monte carlo. While these backgrounds are included in the overall estimates above, the conversion backgrounds are interesting in that the effectiveness of various cuts prescribed for these backgrounds can fail, as will be seen below.

Finally, the analysis for $K_L \rightarrow \pi^0 e^+ e^- \gamma$ benefits from a very useful coincidence with

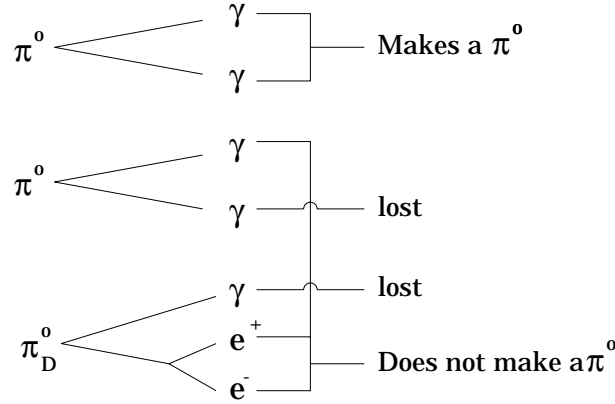


Figure 74. An example of $K_L \rightarrow 3\pi^0_D$ background to $K_L \rightarrow \pi^0 e^+ e^- \gamma$. In this case, a photon was lost from the Dalitz π^0 decay and a photon was lost from one of the normal π^0 decays. Since an intact π^0 survives, but the Dalitz π^0 decay is fragmentary, neither the cut on π^0 mass nor the cut against π^0 mass in the Dalitz case are effective here. Rather, cuts are made in the decay kinematics and in the RC counters to reduce this background.

regard to the estimation of systematic error. $K_L \rightarrow 2\pi^0_D$ is the principal background to $K_L \rightarrow \pi^0 e^+ e^- \gamma$ and is removed mostly by a single cut on the $e^+ e^- \gamma$ invariant mass. Thus, like a light switch, we can turn this background on and off at will. We can thus study the efficiency of many of the quantities on which we cut by “turning on” the background and comparing cut efficiencies in data and $K_L \rightarrow 2\pi^0_D$ monte carlo. $K_L \rightarrow 2\pi^0_D$ is also the normalization mode for $K_L \rightarrow \pi^0 e^+ e^- \gamma$.

8.1 Analysis Cuts for $K_L \rightarrow \pi^0 e^+ e^- \gamma$

In order to reduce systematic effects in the ratio of normalization mode to signal mode acceptances, cuts applied in the normalization mode analysis were also applied in the signal mode analysis where feasible. The only cut in the normalization mode analysis that is not used in this signal mode analysis is the requirement that the invariant mass of the $e^+ e^-$ pair and the lone γ match the π^0 mass. Also, the signal mode analysis includes more cuts against backgrounds whose efficiencies can be studied in the normalization mode.

The photons are paired to identify a π^0 as in the normalization mode. Once a pairing

	π^0 Mean	π^0 Resolution	π_D^0 Mean	π_D^0 Resolution
$K_L \rightarrow 2\pi_D^0$				
Winter Data	135.62 ± 0.01	2.42 ± 0.03	134.76 ± 0.01	1.88 ± 0.02
Winter MC	135.42 ± 0.01	2.43 ± 0.02	134.84 ± 0.01	1.87 ± 0.01
Summer Data	135.60 ± 0.01	2.40 ± 0.03	134.72 ± 0.01	1.90 ± 0.02
Summer MC	135.42 ± 0.01	2.42 ± 0.02	134.84 ± 0.01	1.84 ± 0.01
$K_L \rightarrow \pi^0 e^+ e^- \gamma$				
Winter MC	135.3 ± 0.01	1.75 ± 0.01	N/A	N/A
Summer MC	135.3 ± 0.01	1.74 ± 0.01	N/A	N/A

Table 21. Mass resolutions for the various π^0 masses calculated in the normalization mode analysis and the $K_L \rightarrow \pi^0 e^+ e^- \gamma$ analysis. The mass peaks in each case were fit to a single Gaussian in a close range about each peak. In all cases, the fit χ^2 was between 1.5 and 2.0, and the quoted errors are the fit errors. The quoted resolution is equal to the sigma of the fit Gaussian.

choice is made, a *neutral vertex* can be calculated also as in the normalization mode. The invariant mass of the entire $ee\gamma\gamma\gamma$ final state is then calculated using the neutral vertex. Since the calculation necessarily forces the mass of the photon pair to be the nominal π^0 mass, the mass resolution of the entire final state is improved.

8.1.1 Cuts to Identify $K_L \rightarrow \pi^0 e^+ e^- \gamma$

The basic cuts to identify $K_L \rightarrow \pi^0 e^+ e^- \gamma$ are

- $130 < m_{\gamma\gamma} < 140 \text{ MeV}/c^2$,
- $490 < m_{e^+e^-\gamma\gamma\gamma} < 505 \text{ MeV}/c^2$,
- $P_t^2 < 300 \text{ MeV}^2/c^2$.

The π^0 mass cut is a 3 sigma cut, as in the normalization mode analysis. The $e^+e^-\gamma\gamma\gamma$ mass cut is only 2.3 sigma however. The reason why these cuts are slightly tighter in the $K_L \rightarrow \pi^0 e^+ e^- \gamma$ analysis reflects the effort to isolate a clean signal.

	$e^+e^-\gamma\gamma\gamma$ Mean	$e^+e^-\gamma\gamma\gamma$ Resolution
$K_L \rightarrow 2\pi_D^0$		
Winter Data	497.57 ± 0.02	3.41 ± 0.02
Winter MC	497.42 ± 0.02	3.37 ± 0.02
Summer Data	497.76 ± 0.02	3.44 ± 0.03
Summer MC	497.39 ± 0.02	3.38 ± 0.02
$K_L \rightarrow \pi^0 e^+ e^- \gamma$		
Winter MC	497.5 ± 0.01	3.05 ± 0.02
Summer MC	497.5 ± 0.01	3.04 ± 0.02

Table 22. Mass resolutions for $e^+e^-\gamma\gamma\gamma$ mass calculated in the normalization mode analysis and the $K_L \rightarrow \pi^0 e^+ e^- \gamma$ analysis. The mass peaks in each case were fit to a single Gaussian in a close range about each peak. In all cases, the fit χ^2 was between 1.5 and 2.0, and the quoted errors are the fit errors. The quoted resolution is equal to the sigma of the fit Gaussian.

8.1.2 Cuts to Remove Backgrounds

8.1.2.1 $K_L \rightarrow 2\pi_D^0$

The monkey with the largest and most obvious wrench is $K_L \rightarrow 2\pi_D^0$. In order to study this background, we plot the invariant mass of the e^+e^- pair and lone γ , shown in figure 75. In the figure, it can be seen that the distribution of events is symmetrical about the π^0 mass with a low side tail 3.5 orders of magnitude down from the peak. The presence of the low side tail is due to final state radiation of the e^+e^- pair, soft bremsstrahlung in the vacuum window or in the TRD radiator material. In any case, the shape is well predicted by the monte carlo except for a small shift in the mean. The analysis requires that the $e^+e^-\gamma$ mass will not lie within the range $115 < m_{e^+e^-\gamma} < 145$ MeV/ c^2 .

However, this cut does not remove all of the $K_L \rightarrow 2\pi_D^0$ background. We need to consider cases where we have chosen the wrong pairing of photons. These cases are illustrated in figure 76. After making all analysis cuts except for cuts against $K_L \rightarrow 2\pi_D^0$ in the wrong pairings, figure 77 shows the distribution of $M_{ee\gamma}$ for the 2nd best pairing and




Figure 75. $M_{ee\gamma}$, the invariant mass of the electron-positron pair and the lone photon, is plotted for $K_L \rightarrow 2\pi_D^0$ monte carlo and E799 data. The upper plot shows the $M_{ee\gamma}$ distribution for monte carlo events that satisfy the $K_L \rightarrow \pi^0 e^+ e^- \gamma$ analysis except for the cut against $M_{ee\gamma}$ in the best pairing hypothesis. While a cut at 145 MeV/c^2 is sufficient on the high side, a cut must be placed at 115 MeV/c^2 on the low side to eliminate $K_L \rightarrow 2\pi_D^0$ events with final state radiation, which comprise the low-side tail about 3.5 orders of magnitude down from the peak. The lower plot shows the same $M_{ee\gamma}$ distribution for E799 data satisfying the $K_L \rightarrow \pi^0 e^+ e^- \gamma$ analysis except for the cut against $M_{ee\gamma}$ in the best pairing hypothesis with $K_L \rightarrow 2\pi_D^0$ monte carlo overlaid. The normalization is absolute, and one can see the good agreement between data and monte carlo in the low side tail.




Figure 76. The algorithm used to identify π^0 decay products in this analysis is a simple one. There are three possible ways to make pairs of two photons given three photons. The pairing hypothesis with the invariant mass closest to the nominal π^0 mass is chosen as the best pairing. This calculation is done with 4-vectors assuming that the photons originate at the charged vertex. Occasionally, the “best” pairing hypothesis is not the correct pairing. In this case, the lone photon and e^+e^- pair may appear to have an invariant mass different from the π^0 mass.

figure 78 shows the distribution of $M_{ee\gamma}$ for the worst pairing possibility. Since the mis-paired backgrounds from $K_L \rightarrow 2\pi_D^0$ are much smaller than the correctly paired background, the final state radiation is of little concern, and we can make symmetric cuts about the π^0 mass. The $e^+e^-\gamma$ mass in the 2nd best pairing possibility is required to lie outside of the range $125 < M_{ee\gamma} < 145 \text{ MeV}/c^2$, while in the worst pairing, it is required to lie outside of the range $130 < M_{ee\gamma} < 140 \text{ MeV}/c^2$. A systematic error can be associated with this cut using the normalization mode. In the normalization mode, an energy shift of $0.2 \text{ MeV}/c^2$ in the mass scale exists between data and monte carlo with the data higher. (The maximum shift observed in table 21 $K_L \rightarrow 2\pi_D^0 \pi^0$ masses is used.) Applying this energy shift to the $K_L \rightarrow \pi^0 e^+ e^- \gamma$ monte carlo by varying the mass cut intervals upwards by $0.2 \text{ MeV}/c^2$ gives an estimate of 0.2% for the systematic error.

After these cuts are made, the efficiency for a $K_L \rightarrow 2\pi_D^0$ event to fake a $K_L \rightarrow \pi^0 e^+ e^- \gamma$ event is $(4.0 \pm 1.2) \times 10^{-7}$. This number represents the full monte carlo calculation of the efficiency including conversions and bremsstrahlung in the detector material upstream of the spectrometer, and accidentals. The number of $K_L \rightarrow 2\pi_D^0$ events expected in the final $K_L \rightarrow \pi^0 e^+ e^- \gamma$ sample is (2.0 ± 0.8) based on a monte carlo sample equivalent to 2.5 times the entire KTeV dataset. The error in this background is statistical only and will be included in the final estimate of systematic error for $\text{BR}(K_L \rightarrow \pi^0 e^+ e^- \gamma)$.

To summarize, the cuts designed to get rid of $K_L \rightarrow 2\pi_D^0$ backgrounds are




Figure 77. The difference of $M_{ee\gamma}$ and M_{π^0} is plotted for $K_L \rightarrow 2\pi_D^0$ monte carlo and E799 data for the 2nd best pairing hypothesis. The upper plot shows the $|M_{ee\gamma} - M_{\pi^0}|$ distribution for monte carlo events that satisfy the $K_L \rightarrow \pi^0 e^+ e^- \gamma$ analysis except for the cut against $M_{ee\gamma}$ in the 2nd best pairing hypothesis. Since there is not as much of this background as in the best pairing case (see figure 75,) we can treat the $K_L \rightarrow 2\pi_D^0$ background peak as symmetric in the 2nd best pairing case. A cut at ± 10 MeV/c² relative to the nominal π^0 mass is made to eliminate this mis-pairing background. The lower plot shows the same distribution for E799 data satisfying the $K_L \rightarrow \pi^0 e^+ e^- \gamma$ analysis except for the cut against $M_{ee\gamma}$ in the 2nd best pairing hypothesis with $K_L \rightarrow 2\pi_D^0$ monte carlo overlaid.




Figure 78. The difference of $M_{ee\gamma}$ and M_{π^0} is plotted for $K_L \rightarrow 2\pi_D^0$ monte carlo and E799 data. The upper plot shows the $|M_{ee\gamma} - M_{\pi^0}|$ distribution for monte carlo events that satisfy the $K_L \rightarrow \pi^0 e^+ e^- \gamma$ analysis except for the cut against $M_{ee\gamma}$ in the worst pairing hypothesis. Since there is not as much of this background as in the best pairing case (see figure 75,) we can treat the $K_L \rightarrow 2\pi_D^0$ background peak as symmetric in the 2nd best pairing case. A cut at ± 5 MeV/c² relative to the nominal π^0 mass is made to eliminate this mis-pairing background. The lower plot shows the same distribution for E799 data satisfying the $K_L \rightarrow \pi^0 e^+ e^- \gamma$ analysis except for the cut against $M_{ee\gamma}$ in the 2nd best pairing hypothesis with $K_L \rightarrow 2\pi_D^0$ monte carlo overlaid.

Background from $3\pi^0$ Dalitz			
Final State			Cuts Against
π^0	π^0	π^0_D	
$\gamma \cancel{\gamma}$	$\gamma \cancel{\gamma}$	$e^+e^- \gamma$	$130 < M_{\gamma\gamma} < 140 \text{ MeV}/c^2$ $M_{ee\gamma} < 115 \text{ MeV}/c^2$ and $> 150 \text{ MeV}/c^2$
$\cancel{\gamma}\cancel{\gamma}$	$\gamma \gamma$	$e^+e^- \gamma$	$M_{ee\gamma} < 115 \text{ MeV}/c^2$ and $> 150 \text{ MeV}/c^2$
$\gamma \cancel{\gamma}$	$\gamma \gamma$	$e^+e^- \cancel{\gamma}$	Kinematics
$e^+e^- \gamma \gamma \gamma$ with fusions			Shape chi-square

Figure 79. A table of simple ways $K_L \rightarrow 3\pi^0_D$ can contribute to $K_L \rightarrow \pi^0 e^+ e^- \gamma$ backgrounds. In addition to the cases listed above, more particles can be lost or fused and replaced by accidentals. However, such cases are more rare.

- Cut if $115 < M_{ee\gamma} < 145 \text{ MeV}/c^2$ in the best pairing hypothesis,
- Cut if $125 < M_{ee\gamma} < 145 \text{ MeV}/c^2$ in the second best pairing hypothesis,
- Cut if $130 < M_{ee\gamma} < 140 \text{ MeV}/c^2$ in the worst pairing hypothesis.

8.1.2.2 $K_L \rightarrow 3\pi^0_D$

The monkey with the second largest wrench is $K_L \rightarrow 3\pi^0_D$. This background is at the same time interesting and difficult to analyze because it is a fragmentary background in which different kinds of fragments can pose different kinds of problems. In the minimal case, at least two photons must be lost from the final state either through the beam holes, outside of the fiducial volume of the detector, or in fusions in the CsI calorimeter. The different ways $K_L \rightarrow 3\pi^0_D$ can pose a background threat with a minimal loss of photons and the cuts against them are enumerated in table 79, and an illustration is in figure 74. This table does not tell the whole story, however; three photons may be lost and one replaced by an accidental. In this case, mass cuts may not be effective, and kinematical cuts and photon veto cuts will need to be made. Fortunately, however, these occurrences are more rare because one has to lose yet another photon as well as pick up an accidental

of the right energy to fake a π^0 . The minimal cases described in the table are therefore more probable and will therefore be the focus of our cuts.

Let us dispose of the $K_L \rightarrow 3\pi_D^0$ background piecemeal. In the case where one or more photon fusions occur in the CsI calorimeter, the background is suppressed by a cut against the cluster shape χ^2 variable. The cluster shape χ^2 is calculated by comparing the energy in each crystal of the cluster to an expected energy calculated from GEANT generated shower shapes. The distributions of cluster shape χ^2 for photons from $K_L \rightarrow 2\pi_D^0$ is shown in figure 80. A cut at a shape χ^2 of 5.0 is sufficient to remove background from fusions. From studying the efficiency of this cut in the normalization sample, we assign a systematic error of 1.8% due to this cut.

To do this, the background $K_L \rightarrow 2\pi_D^0$ sample was “turned on” by removing the cut against $e^+e^-\gamma$ invariant mass in the best pairing hypothesis. The efficiency of the $K_L \rightarrow 2\pi_D^0$ sample to pass this cut in the data was compared to the efficiency in the monte carlo giving a difference of 0.6% per photon. Since there are three photons, and since the data/monte carlo difference in the shape χ^2 variable is at worst a scale change, the per photon figure was multiplied by three to get the above error.

We must have a special cut against this possibility because, unlike the other cases, backgrounds from pure fusions involve no loss of energy from the event so that a final mass cut on $e^+e^-\gamma\gamma\gamma$ invariant mass will not remove them. Even so, in the case of $K_L \rightarrow 3\pi_D^0$, a fusion involving two distinct pairs of photons without accidentals cannot satisfy the requirement on the π^0 mass except accidentally.

For the other cases listed in the table, backgrounds where a single photon is missing from each of the $\pi^0(\gamma\gamma)$ in the final state are suppressed because the remaining photons will not tend to make a π^0 so that they tend to fail the requirement that there be two photons that make a π^0 . Backgrounds where an entire $\pi^0(\gamma\gamma)$ is missing are suppressed because the (correctly paired) $e^+e^-\gamma$ make do a π^0 mass and are therefore suppressed by the cuts against $K_L \rightarrow 2\pi_D^0$. Backgrounds where a soft photon is missing from the $e^+e^-\gamma$ and a soft photon is missing from one of the $\pi^0(\gamma\gamma)$ present the most serious potential problem.

In order to remove remaining $K_L \rightarrow 3\pi_D^0$ background, a cut is made against the

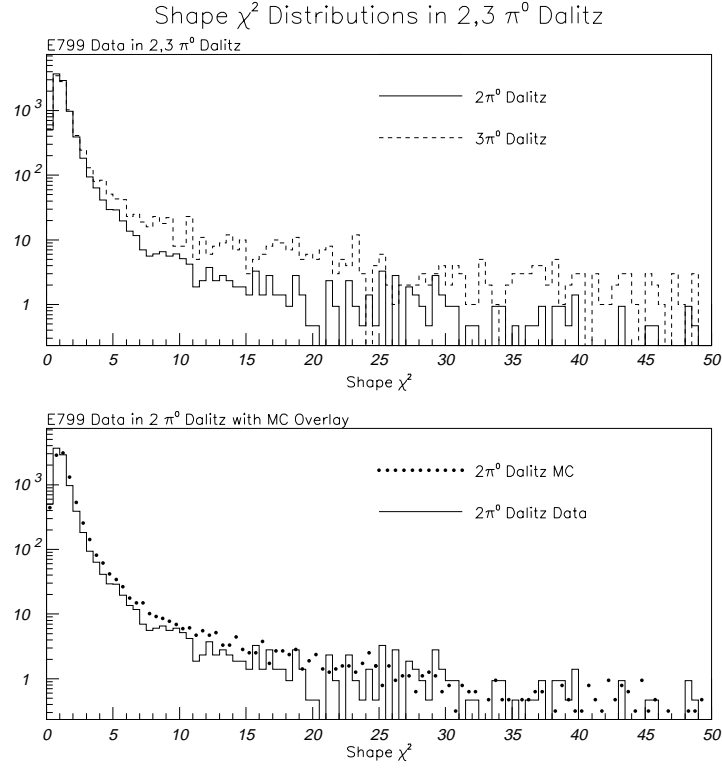


Figure 80. The distribution of CsI cluster shape chi-square for photons is plotted for $K_L \rightarrow 2\pi_D^0$ and $K_L \rightarrow 3\pi_D^0$ events. The upper plot shows the distribution of CsI cluster shape chi-square for $K_L \rightarrow 2\pi_D^0$ events that satisfy the $K_L \rightarrow \pi^0 e^+ e^- \gamma$ analysis except for the cut against $M_{ee\gamma}$ invariant mass in the best pairing hypothesis. Overlaid is a sample of $K_L \rightarrow 3\pi_D^0$ monte carlo events that satisfy the $K_L \rightarrow \pi^0 e^+ e^- \gamma$ analysis except for the cut against $K_L \rightarrow 3\pi_D^0$ kinematics and the final kaon mass cut. The excess in the case of $K_L \rightarrow 3\pi_D^0$ is due to the existence of overlapping photon clusters in the CsI. Since $K_L \rightarrow 3\pi_D^0$ with two extremely close fusions may pass mass and pt cuts, it is a serious background problem. A cut is made at a chi-square value of 5 to reject fusions. The lower plot shows the cluster shape chi-square variable for photons in a similarly defined $K_L \rightarrow 2\pi_D^0$ sample in the E799 data with $K_L \rightarrow 2\pi_D^0$ monte carlo overlay. The agreement between data and monte carlo is good.

kinematics of the decay. The general details of this kinematical cut are presented in Appendix C. To make this cut, the final state of the event is broken into two groups. The first group consists of the e^+e^- pair and the second group consists of the $\gamma\gamma\gamma$. Then, the square of the momentum of the e^+e^- system in the kaon rest frame along the kaon boost direction is plotted against the $\gamma\gamma\gamma$ invariant mass. A cut can be made in this plane that in general should distinguish fragmentary backgrounds from intact final states. In order to generate an efficiently contoured cut in this plane, a cut is made in the variable κ , defined by

$$\kappa = \log P_{ee}^{\parallel*} - 15.6M_{3\gamma} + 4.6, \quad (8.1)$$

where $P_{ee}^{\parallel*}$ is the component of kaon CM frame momentum of the e^+e^- system parallel to the kaon direction in the lab frame. With κ defined this way, non-fragmentary $K_L \rightarrow \pi^0 e^+ e^- \gamma$ events should roughly satisfy $\kappa = 0$. Note also that since $K_L \rightarrow 2\pi_D^0$ events are also non-fragmentary, they also populate the region around $\kappa = 0$. We will take advantage of this fact to estimate a systematic error in this cut as in the case of the shape χ^2 variable. Figure 81 shows the distribution in κ for $K_L \rightarrow \pi^0 e^+ e^- \gamma$ and surviving $K_L \rightarrow 3\pi_D^0$ events after all cuts except for a mass cut on the final state. This analysis requires $\kappa < 0.5$. From studying the agreement between data and monte carlo in the normalization mode (see figure 82,) we get an estimate of systematic error of 0.6%. We have obtained this estimate by “turning on” the $K_L \rightarrow 2\pi_D^0$ background by removing the cut against $e^+e^- \gamma$ invariant mass in the best pairing hypothesis. The efficiency of the $K_L \rightarrow 2\pi_D^0$ sample to pass the cut on κ in the data was compared to the efficiency to pass the cut in the $K_L \rightarrow 2\pi_D^0$ monte carlo. The ratio of efficiencies amount to a 0.6% difference. This estimate can be applied to the $K_L \rightarrow \pi^0 e^+ e^- \gamma$ analysis on the strength of the close agreement between $K_L \rightarrow 2\pi_D^0$ monte carlo and $K_L \rightarrow \pi^0 e^+ e^- \gamma$ monte carlo in the region of the cut as shown in the bottom of figure 82.

A cut on photon vetoes beyond the trigger level was also considered at this point. The photon vetoes can be used to reject background from $K_L \rightarrow 3\pi_D^0$ where one or more photons were lost in general. In order to study the photon vetoes, two special data samples were prepared. First, a sample of $K_L \rightarrow 2\pi_D^0$ could be isolated by relaxing the cut against the $e^+e^- \gamma$ mass matching a π^0 . Second, a $K_L \rightarrow 3\pi_D^0$ sample could be isolated

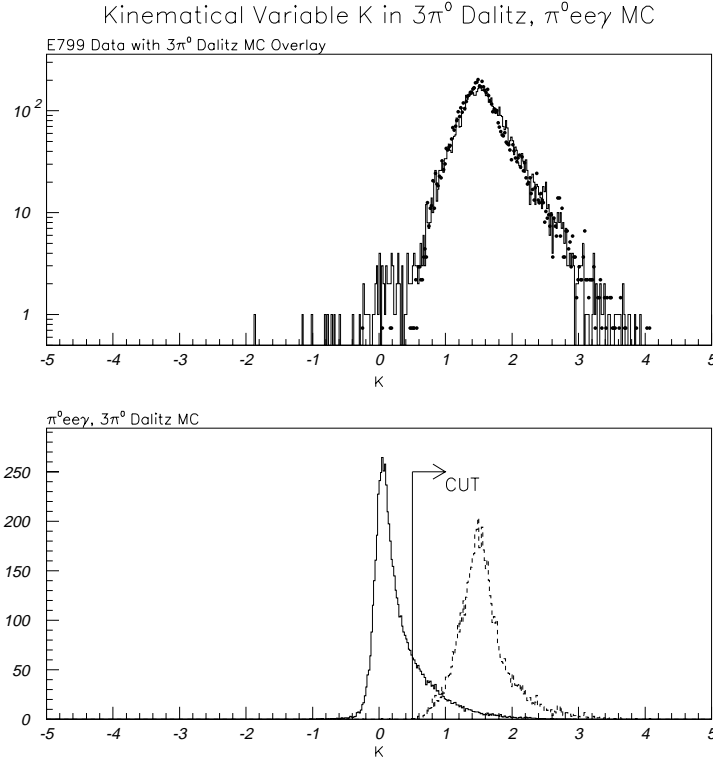


Figure 81. Kinematic variable κ , defined in appendix D, for $K_L \rightarrow 3\pi_D^0$ rejection. The top plot shows E799 data in the $K_L \rightarrow \pi^0 e^+ e^- \gamma$ sample after all cuts (except this one and the cut on $e^+ e^- \gamma \gamma \gamma$ mass) with a $K_L \rightarrow 3\pi_D^0$ monte carlo overlay. The normalization is absolute, and as expected, the background at this stage is dominated by $K_L \rightarrow 3\pi_D^0$. The bottom plot shows the $K_L \rightarrow \pi^0 e^+ e^- \gamma$ monte carlo distribution in κ also with $K_L \rightarrow 3\pi_D^0$ monte carlo. The data at $\kappa = 0$ are populated by non-fragmentary $K_L \rightarrow \pi^0 e^+ e^- \gamma$ events. The cut effectively removes the threat of $K_L \rightarrow 3\pi_D^0$ background.

by relaxing the cut on κ . Figure 83 shows the maximum signal deposit in any SA counter for each of these samples with monte carlo overlay, and figure 84 shows the maximum signal deposit in any RC. Clearly, it is much more efficient to cut on the RCs. Therefore, a cut was placed requiring the maximum RC signal to be less than 0.5 MiP. This is about 30% more stringent than the trigger requirement. A systematic error of 0.5% is estimated from the $K_L \rightarrow 2\pi_D^0$ data comparison to monte carlo. We have obtained this estimate by “turning on” the $K_L \rightarrow 2\pi_D^0$ background by removing the cut against $e^+ e^- \gamma$ invariant mass in the best pairing hypothesis. The efficiency of the $K_L \rightarrow 2\pi_D^0$ sample to pass




Figure 82. Kinematic variable κ , defined in appendix D, is plotted for E799 data with the cut against $K_L \rightarrow 2\pi_D^0$ removed with a $K_L \rightarrow 2\pi_D^0$ monte carlo overlay (top). The normalization is absolute, and the $K_L \rightarrow 2\pi_D^0$ shape is well predicted by the monte carlo. The bottom plot compares the distributions in κ for $K_L \rightarrow 2\pi_D^0$ and $K_L \rightarrow \pi^0 e^+ e^- \gamma$ monte carlo.

the cut on RC energy in the data was compared to the efficiency to pass the cut in the $K_L \rightarrow 2\pi_D^0$ monte carlo. The ratio of efficiencies amount to a 0.5% difference.

After all cuts are made, the efficiency for a $K_L \rightarrow 3\pi_D^0$ event to fake a $K_L \rightarrow \pi^0 e^+ e^- \gamma$ event is measured to be $(5 \pm 3.5) \times 10^{-10}$. This number represents the full monte carlo calculation of the efficiency including conversions and bremsstrahlung in the detector material upstream of the spectrometer, lost particles, and accidentals. The number of $K_L \rightarrow 3\pi_D^0$ events expected in the background is estimated to be (1.0 ± 0.7) events. The large systematic error on this number is due to the statistics in the monte carlo sample. Unlike $K_L \rightarrow 2\pi_D^0$, it was difficult to make a monte carlo sample as big as the KTeV dataset. This is because the higher branching ratio ($\text{BR}(K_L \rightarrow 3\pi_D^0)$) is about

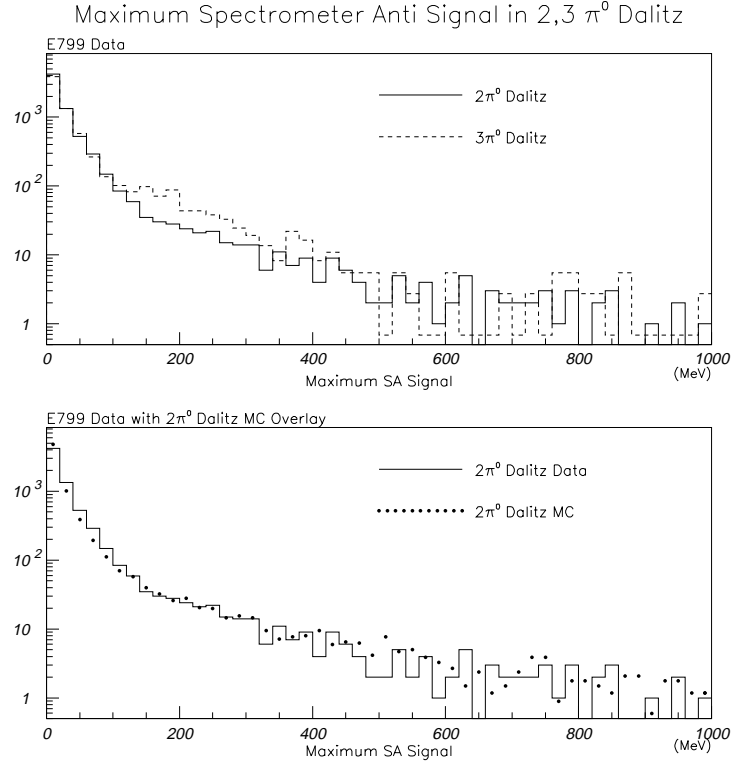


Figure 83. The distribution of maximum energy deposited in a spectrometer anti (SA) is plotted for $K_L \rightarrow 2\pi_D^0$ and $K_L \rightarrow 3\pi_D^0$ events. Energy is summed in each SA counter, and the maximum energy deposit over the set of SA counters is used to make these plots. The upper plot shows the distribution of maximum energy deposit in the SAs for $K_L \rightarrow 2\pi_D^0$ events that satisfy the $K_L \rightarrow \pi^0 e^+ e^- \gamma$ analysis except for the cut against $M_{ee\gamma}$ invariant mass in the best pairing hypothesis. Overlaid is a sample of $K_L \rightarrow 3\pi_D^0$ monte carlo events that satisfy the $K_L \rightarrow \pi^0 e^+ e^- \gamma$ analysis except for the cut against $K_L \rightarrow 3\pi_D^0$ kinematics and the final kaon mass cut. The $K_L \rightarrow 3\pi_D^0$ show a slight enhancement from lost photons striking an SA, while the $K_L \rightarrow 2\pi_D^0$ sample represents a null case for comparison. The lower plot shows a similarly defined set of $K_L \rightarrow 2\pi_D^0$ events in the E799 data with $K_L \rightarrow 2\pi_D^0$ monte carlo overlay. The agreement between data and monte carlo is excellent. However, the distinction between $K_L \rightarrow 3\pi_D^0$ and the null case is not so great. Therefore, we do not cut on the SAs.

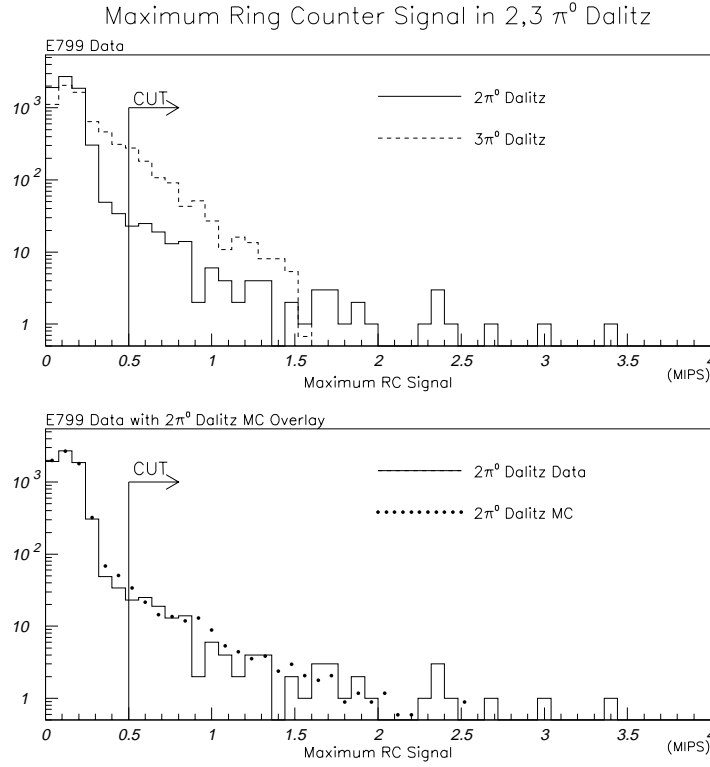


Figure 84. The distribution of maximum energy deposited in a ring counter (RC) is plotted for $K_L \rightarrow 2\pi_D^0$ and $K_L \rightarrow 3\pi_D^0$ events. Energy is summed in each RC counter, and the maximum energy deposit over the set of RC counters is used to make these plots. The upper plot shows the distribution of maximum energy deposit in the RCs for $K_L \rightarrow 2\pi_D^0$ events that satisfy the $K_L \rightarrow \pi^0 e^+ e^- \gamma$ analysis except for the cut against $M_{ee\gamma}$ invariant mass in the best pairing hypothesis. Overlaid is a sample of $K_L \rightarrow 3\pi_D^0$ monte carlo events that satisfy the $K_L \rightarrow \pi^0 e^+ e^- \gamma$ analysis except for the cut against $K_L \rightarrow 3\pi_D^0$ kinematics and the final kaon mass cut. The $K_L \rightarrow 3\pi_D^0$ show a slight enhancement from lost photons striking an RC, while the $K_L \rightarrow 2\pi_D^0$ sample represents a null case for comparison. The lower plot shows a similarly defined set of $K_L \rightarrow 2\pi_D^0$ events in the E799 data with $K_L \rightarrow 2\pi_D^0$ monte carlo overlay. The agreement between data and monte carlo is excellent. A cut is at 0.5 MIPS (equivalent signal *Minimum Ionizing Particle*) to reject $K_L \rightarrow 3\pi_D^0$ background.




Figure 85. The E799 dataset is plotted with $K_L \rightarrow 3\pi_D^0$ and $K_L \rightarrow 2\pi_D^0$ monte carlo overlay after all cuts except for the cut against the kinematical variable, κ . The normalization is absolute.

0.76%) requires the generation of more events. Also, since accidentals in the CsI were a concern, each event had to be fully built before checking the number of hardware clusters. Therefore, a sample of size comparable to only 1.0 times the KTeV dataset was generated.

To summarize, the cuts designed to get rid of $K_L \rightarrow 3\pi_D^0$ backgrounds are

- Require shape $\chi^2 < 5.0$,
- Require kinematical variable $\kappa < 0.5$,
- Require maximum energy in the RCs < 0.5 Mip.

Figure 85 shows the $e^+e^-\gamma\gamma$ invariant mass distribution after all cuts except for the cut against κ . At this stage of analysis, the signal events are clearly seen, and the principle remaining backgrounds are seen to be $K_L \rightarrow 2\pi_D^0$ and $K_L \rightarrow 3\pi_D^0$.




Figure 86. The $K_L \rightarrow \pi^0 \gamma \gamma$ final state. This can cause a background problem for the $K_L \rightarrow \pi^0 e^+ e^- \gamma$ analysis if one of the direct photons converts to an $e^+ e^-$ pair in the vacuum window. This case is suppressed by the requirement that the decay vertex occurs upstream of the vacuum window.

8.1.2.3 Other Backgrounds

A potentially serious background could come from $K_L \rightarrow \pi^0 \gamma \gamma$. However, we are fortunate that the rate in this mode is already small [Alavi 99b].

$$BR(K_L \rightarrow \pi^0 \gamma \gamma) = 1.68 \times 10^{-6} \quad (8.2)$$

First, we must worry about the case where the π^0 decays to $e^+ e^- \gamma$. This case is further suppressed by the branching ratio; and also there are no two photons in the final state that naturally make a π^0 . In the case that the free $\gamma \gamma$ part of the $K_L \rightarrow \pi^0 \gamma \gamma$ decay make a π^0 mass, already a small possibility from the differential decay spectrum in $m_{\gamma \gamma}$ the remaining $e^+ e^- \gamma$ part of the final state makes a π^0 mass. Therefore this possibility is already suppressed by the analysis cuts. Another possibility is that a photon from the $\pi^0(\gamma \gamma)$ decay part of $K_L \rightarrow \pi^0 \gamma \gamma$ will convert to an $e^+ e^-$ pair in the vacuum window or material upstream of the spectrometer. This case is suppressed for the same reasons as above. The last case, where one of the photons from the direct $\gamma \gamma$ part of the $K_L \rightarrow \pi^0 \gamma \gamma$ final state converts in the window, is potentially more serious. If a good vertex is obtained, then the decay final state can fake $K_L \rightarrow \pi^0 e^+ e^- \gamma$ and the $e^+ e^- \gamma$ invariant mass spectrum will resemble that in $K_L \rightarrow \pi^0 e^+ e^- \gamma$. However, in the event that a good vertex is obtained, then the event will be suppressed by the requirement that the vertex Z position

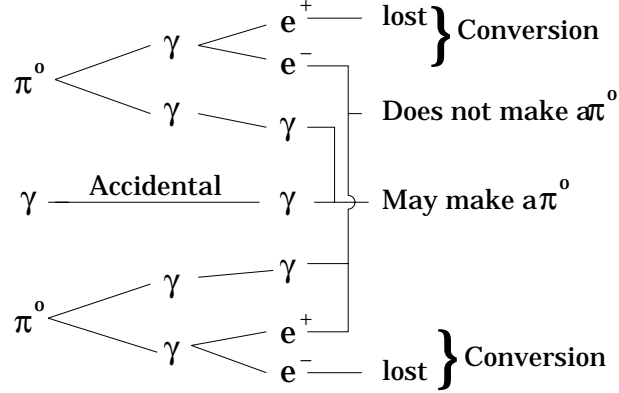


Figure 87. An example of a double conversion event from $K_L \rightarrow 2\pi^0$ where two photons convert in the vacuum window. This can in principle contribute to $K_L \rightarrow \pi^0 e^+ e^- \gamma$ background.

is upstream of the vacuum window. The conversion probability is also small, so that in the final analysis, we can get away without inventing any special cuts against it.

The tale of backgrounds is not yet full with respect to $K_L \rightarrow 2\pi^0$ and $K_L \rightarrow 3\pi^0$. These backgrounds can also contribute to background through the mechanisms of window conversion. Although the conversion probability per photon is about 1%, the total rate of this background can be higher than the similar Dalitz backgrounds because of the multiplicity of photons. However, in the case of a reconstructed $K_L \rightarrow 2\pi^0$ or $K_L \rightarrow 3\pi^0$ single conversion event, the charged vertex reconstructs at the vacuum window and is therefore also suppressed by the requirement that the charged vertex lies upstream of the vacuum window.

In the case where two photons convert and two soft tracks are swept out of the detector at the analysis magnet, we can have a serious background problem. In this very interesting scenario, depicted in figure 88, two soft and oppositely charged electrons, one from each conversion pair, may be swept out of the detector by the analysis magnet. This has been modeled special monte carlo samples and can be shown to be very small.

Natural four track backgrounds, which come primarily from $K_L \rightarrow 3\pi^0, 2\pi^0 \rightarrow e^+ e^- \gamma$ “double Dalitz” decays, can be a problem if only one photon is lost, but an oppositely charged pair of electron tracks is swept out of the detector at the analysis magnet. Although suppressed by the presence of a second Dalitz decay, this background has presented a serious problem for other rare decay searches in KTeV. Natural four track background




Figure 88. A double conversion event in the charged spectrometer. In this scenario, two photons from a $K_L \rightarrow 2\pi^0$ or $K_L \rightarrow 3\pi^0$ event convert in the window. The opening angle between the tracks of each conversion pair is small, so that each pair looks like a single track in the upstream chambers. In the unlikely event that an oppositely charged pair of soft electrons is swept out by the analysis magnet, background can result since the resulting tracks may appear to make a good vertex upstream of the vacuum window. This background has been modeled and is very small, however, and does not contribute to the $K_L \rightarrow \pi^0 e^+ e^- \gamma$ sample at this level.

is distinguished by the presence of extra in-time hit pairs in DC1 and DC2, the chambers upstream of the analysis magnet. Therefore, to eliminate the contribution of this background, events are which have extra in-time hit pairs in DC2.¹ Note that while this cut is efficacious against genuine physics four track backgrounds, it is not effective against the double conversion backgrounds. Vacuum window conversions are distinguished by very low $e^+ e^-$ invariant mass pairs. A conversion pair by contrast often mimics single tracks in the upstream chambers so that double conversion events are not accompanied by extra in-time hit pairs in the upstream chambers.

To summarize, the backgrounds mentioned in this subsection contribute very little to the $K_L \rightarrow \pi^0 e^+ e^- \gamma$ sample. A cut against extra in-time hit pairs in DC2 has been

¹Making the cut in either chamber is sufficient to blow away the natural four-track backgrounds. Cutting on both chambers is unnecessarily costly in efficiency, while DC2 is in general less noisy than DC1 and is therefore better suited to make this cut.




Figure 89. A natural four track event, such as $K_L \rightarrow 3\pi^0$ double Dalitz, has tracks that are well separated in the upstream chambers. Therefore, if two oppositely charged tracks get swept away at the analysis magnet, the background can be suppressed by making a cut against the possibility of in-time hit pairs in the upstream chambers.

added to guard against the possibility of natural four-track backgrounds contributing. The integrated number of events estimated from these sources is less than 0.1 event, and will not be considered further.

8.1.2.4 A Monkey With Two Wrenches : $K_L \rightarrow e^+e^-\gamma$

The decay $K_L \rightarrow e^+e^-\gamma$ can be a background to $K_L \rightarrow \pi^0 e^+e^-\gamma$ through the internal bremsstrahlung of one or both of its electrons. First, $K_L \rightarrow e^+e^-\gamma\gamma$ if accompanied by an accidental photon or a bremsstrahlung photon from the vacuum window can contribute. Since $K_L \rightarrow e^+e^-\gamma\gamma$ is related to the decay $K_L \rightarrow e^+e^-\gamma$ in which one of the photons comes from internal bremsstrahlung so that in either case, we expect to find that one of the photons is nearly collinear with an electron track. In order to get a handle on this background, we look at the smallest distance of a hardware cluster in the CsI to the extrapolated position of an upstream track segment of one of the electrons, D_{min} . Though this background is expected to be small in the data from monte carlo studies, a cut




Figure 90. For each event, we can find the minimum distance between the extrapolated position of an upstream segment of a charged track at the CsI and a hardware cluster in the CsI. This variable is plotted above for $K_L \rightarrow \pi^0 e^+ e^- \gamma$ monte carlo and for remaining events in the E799 data. The significance of this variable is that bremsstrahlung photons are distinguished by their closeness to the associated charged track. (Since charged tracks are bent by the analysis magnet, we use extrapolated upstream track segments for this calculation.) One can see a small peak of 9 events in the data that seem to correspond to bremsstrahlung events and are not explained by the $K_L \rightarrow \pi^0 e^+ e^- \gamma$ monte carlo.

requiring $D_{min} > 1.25$ cm (0.5 of a small crystal width) can reduce this type of background by 95% for only a very small signal inefficiency of 0.5%. Without the above cut, we expect that the $K_L \rightarrow e^+ e^- \gamma \gamma$ background with accidentals and window bremsstrahlung would contribute about 1.0 event to the $K_L \rightarrow \pi^0 e^+ e^- \gamma$ sample in the region $D_{min} < 1.25$ cm. In fact, there are nine such events. Figure 90 shows the distribution of D_{min} for the $K_L \rightarrow \pi^0 e^+ e^- \gamma$ monte carlo and for the E799 $K_L \rightarrow \pi^0 e^+ e^- \gamma$ dataset. Since this was completely unexpected, these nine events were investigated further. To do this, I looked




Figure 91. The angle ϕ between the e^+e^- plane and the $\gamma\gamma$ plane is depicted above. To get the angle between the decay planes, the relation $\cos \phi = \frac{(\gamma_1 \times \gamma_2)(e_1 \times e_2)}{norm}$,

at the angles between the normals of the e^+e^- decay plane and the decay plane of each $\gamma\gamma$ pair. This angle is defined for any pair of two of three photons.

$$\cos \Delta\phi = \frac{(\gamma_1 \times \gamma_2)(e_1 \times e_2)}{norm}. \quad (8.3)$$

An example of such an angle is shown in figure 91. Since there are three such angles, there are three corresponding $\Delta\phi$ angles to check. In figure 92, the data are plotted in a scatterplot of two of these $\Delta\phi$ variables. Seven of the nine events that lie in the region $D_{min} < 1.25$ cm also satisfy $\Delta\phi < 0.2$ radians simultaneous in both plotted angles. These seven events also satisfy that condition in the third $\Delta\phi$ variable. This means that seven of the nine events in the $D_{min} < 1.25$ cm region are nearly coplanar. This is inconsistent with the decay dynamics proposed for $K_L \rightarrow \pi^0 e^+ e^- \gamma$, but it is consistent with the hypothesis that the events in question arise from $K_L \rightarrow e^+ e^- \gamma \gamma$ decay, which is $K_L \rightarrow e^+ e^- \gamma$ decay with two internal bremsstrahlung photons. If so,




Figure 92. $\Delta\phi$ is plotted for two pairings of three photons. Seven of the nine events in the $D_{min} < 1.25$ cm region also satisfy $\Delta\phi < 0.2$ radians simultaneously in two of the angular variables (and in the third as well.) The top plot shows the distribution for $K_L \rightarrow \pi^0 e^+ e^- \gamma$ events and the bottom is the data. Therefore, the seven events are coplanar. This is consistent with $K_L \rightarrow e^+ e^- \gamma \gamma$ decay, aka $K_L \rightarrow e^+ e^- \gamma$ with two internal bremsstrahlung photons.

Background Source	Number Events Expected
$K_L \rightarrow 2\pi_D^0$	2.0 ± 0.8
$K_L \rightarrow 3\pi_D^0$	1.0 ± 1.0
$K_L \rightarrow e^+e^-\gamma\gamma$	0.5 ± 0.1

Table 23. The relevant background contributions to the final $K_L \rightarrow \pi^0 e^+ e^- \gamma$ sample. Backgrounds from double conversion events and four track events are very small and are not included.

then this analysis also represents the first observation of $K_L \rightarrow e^+ e^- \gamma \gamma$! No observation of this process exists in [PDG 98] and to my knowledge no calculation has been done. If we consider the ratio of $K_L \rightarrow e^+ e^- \gamma \gamma$ decay to $K_L \rightarrow e^+ e^- \gamma$ decay, we can estimate $\text{BR}(K_L \rightarrow e^+ e^- \gamma \gamma)$ to be roughly 4.9×10^{-8} based on the $K_L \rightarrow e^+ e^- \gamma \gamma$ branching ratio of 6.5×10^{-7} . Since this decay mode is not implemented in the KTeV monte carlo, it is difficult at the time of writing to estimate the acceptance of this mode. In any event, these coplanar events are removed by the cut against $D_{min} < 1.25$ cm. After this cut, the number of remaining events from $K_L \rightarrow e^+ e^- \gamma \gamma$ plus accidentals is (0.5 ± 0.1) events.

The decay $K_L \rightarrow e^+ e^- \gamma$ can also contribute to the background with the addition of bremsstrahlung photons and/or accidentals. In the case where bremsstrahlung photons are present, this background is reduced by the cut on the minimum distance requirement mentioned above. In the case where photons are purely accidental, then we expect that mass and pt cuts will do the work for us.

8.2 Results on $K_L \rightarrow \pi^0 e^+ e^- \gamma$

8.2.1 Summary of Calculated Backgrounds to $K_L \rightarrow \pi^0 e^+ e^- \gamma$

A summary of the expected backgrounds appears in table 23. The background events will be distributed among the summer and winter according to their weights in the overall flux; 41% and 59% respectively.

8.2.2 Acceptance of $K_L \rightarrow \pi^0 e^+ e^- \gamma$

Using the KTeV monte carlo simulation and the package provided by Fabrizio Gabiani, the acceptance of $K_L \rightarrow \pi^0 e^+ e^- \gamma$ was measured to be $(0.646 \pm 0.004)\%$ in the summer and $(0.660 \pm 0.004)\%$ in the winter. After all cuts, there are 17 events in the summer E799 dataset and 28 events in the winter E799 dataset. The branching ratio is calculated to be

$$BR(K_L \rightarrow \pi^0 e^+ e^- \gamma) = (2.42 \pm 0.38_{\text{STAT}}) \times 10^{-8}, \quad (8.4)$$

where the error is statistical on 41.5 events after background subtraction.

8.2.3 Systematic Error

The detector based systematic errors quoted in the flux measurement of chapter 7 are assumed to cancel in the branching ratio measurement. This is true since the branching ratio depends upon the ratio of the acceptances for $K_L \rightarrow \pi^0 e^+ e^- \gamma$ and $K_L \rightarrow 2\pi_D^0$. However, there are a few exceptions. For the cuts that are not common to the signal and normalization mode analyses, detector systematic errors must be counted. Also, for cases in which the systematic error depends upon detector illumination (which is different between the two modes,) the systematic error must be counted. An example of an illumination effect is the systematic due to drift chamber inefficiencies in the beam region, (see figure 63.) In the figure, it can be seen that there is a different inefficiency due to the drift chambers when 0, 1, or 2 tracks traverse the beam regions. A systematic error can be estimated by measuring the efficiencies in the data using the normalization mode; then the total track weighted efficiencies in $K_L \rightarrow 2\pi_D^0$ and $K_L \rightarrow \pi^0 e^+ e^- \gamma$ monte carlo can be compared using the measured efficiencies. The difference is only in the illumination of the DC beam regions between the two modes. The systematic due to illumination effects is estimated at 0.4%. Finally, the error from the branching ratio of $K_L \rightarrow 2\pi_D^0$ must be counted. These systematic errors appear in table 24. Notably, the zslope systematic will not be included. The total systematic error, including detector

	Systematic Error
BR($K_L \rightarrow 2\pi_D^0$)	3.42%
Background MC Statistics	3.0%
Shape χ^2	1.8%
Cut on κ	0.6%
Cut on RC vetoes	0.5%
Illumination Effects	0.4%
Mass Cuts	0.2%
Total	4.8 %

Table 24. Systematic errors in the measurement of $K_L \rightarrow \pi^0 e^+ e^- \gamma$.

effects and the error in the branching ratio and the monte carlo statistics is 4.8% by adding the above values in quadrature.

8.2.4 Tying It All Together

In summary, the cuts in the $K_L \rightarrow \pi^0 e^+ e^- \gamma$ analysis are largely targeted to remove specific backgrounds. There are also cuts common to the signal mode analysis and the normalization mode analysis to reduce systematic differences as much as possible. Finally, there are cuts designed to isolate the signal itself. A cut-flow analysis table appears in table 25.

The normalization mode cuts are described in detail in chapter 7. The basic cuts to identify $K_L \rightarrow \pi^0 e^+ e^- \gamma$ are

- $130 < m_{\gamma\gamma} < 140 \text{MeV}/c^2$,
- $490 < m_{e^+e^-\gamma\gamma} < 505 \text{MeV}/c^2$,
- $P_t^2 < 300 \text{MeV}^2/c^2$,

where the cut on $\gamma\gamma$ mass is included as part of the normalization mode cuts and the cut on the total final state invariant mass is listed as the “Final Mass Cut.”

The cuts to remove background from $K_L \rightarrow 2\pi_D^0$, the “ $2\pi_D^0$ Killer” cuts, are

- Cut if $115 < M_{ee\gamma} < 145$ MeV/ c^2 in the best pairing hypothesis,
- Cut if $125 < M_{ee\gamma} < 145$ MeV/ c^2 in the second best pairing hypothesis,
- Cut if $130 < M_{ee\gamma} < 140$ MeV/ c^2 in the worst pairing hypothesis.

The cuts to remove background from $K_L \rightarrow 3\pi_D^0$ are

- Require maximum energy in the RCs < 0.5 Mip
- Require shape $\chi^2 < 5.0$, and
- Require kinematical variable $\kappa < 0.5$ (“ $3\pi_D^0$ Killer”).

Finally, cuts to remove other backgrounds are

- $D_{min} > 1.25$ cm. (Gets rid of possible $K_L \rightarrow e^+e^-\gamma\gamma\gamma$ events.)
- No extra in-time hits in DC2 (Gets rid of possible four track backgrounds.)

From the table, one can see that the various targeted cuts are having the desired effect on the relevant backgrounds. For example, the “ $2\pi_D^0$ Killer” cuts have a relative efficiency of 87% for signal but only 0.57% for $K_L \rightarrow 2\pi_D^0$ background, leading to an efficiency weighted rejection factor of 152. Also, the three cuts targeted to remove $K_L \rightarrow 3\pi_D^0$ background 68% efficient for signal but only 3.7% efficient for background, leading to an efficiency weighted rejection of 18.2.

8.2.5 First Observation of $K_L \rightarrow \pi^0 e^+ e^- \gamma$

This analysis represents the first observation of the decay $K_L \rightarrow \pi^0 e^+ e^- \gamma$. The branching ratio is measured to be

$$BR(K_L \rightarrow \pi^0 e^+ e^- \gamma) = (2.42 \pm 0.38_{\text{STAT}} \pm 0.11_{\text{SYS}}) \times 10^{-8}. \quad (8.5)$$

Figure 93 shows the $e^+e^-\gamma$ invariant mass distribution after all cuts. The distinctive shape is well predicted by the monte carlo model, based upon the chiral perturbation

Cut	Number in Data	Combined Acceptance	Normalized $K_L \rightarrow 2\pi_D^0$ MC	Normalized $K_L \rightarrow 3\pi_D^0$ MC
Normalization Mode Cuts	60235	1.48%	51557.2	3127.0
P_t^2	43337	1.24%	39642.0	604.0
$2\pi_D^0$ Killer	1344	1.08%	226.0	446.0
No Extra In- Time Hit Pairs	974	1.01%	202.4	353.0
Minimum Extrap Trk-Cls dist	914	1.01%	136.0	351.0
RC Energy	845	0.97%	129.6	301.0
Shape χ^2	254	0.88%	66.0	85.0
$3\pi_D^0$ Killer	87	0.68%	16.8	13.0
Final Mass Cut	45	0.65%	2.0	1.0

Table 25. Cut flow table for the $K_L \rightarrow \pi^0 e^+ e^- \gamma$ analysis. The normalization mode cuts are discussed in chapter 7. The “ $2\pi_D^0$ Killer” refers to the mass cuts used against $K_L \rightarrow 2\pi_D^0$ in all photon pairing hypotheses. The “ $3\pi_D^0$ Killer” refers to the kinematical cut against $K_L \rightarrow 3\pi_D^0$. The cut values are described in the text. The first column lists the number of events surviving in the data. The second column lists the acceptance (combined winter and summer numbers.) The third and final columns detail the normalized expected backgrounds from the two major background sources, $K_L \rightarrow 2\pi_D^0$ and $K_L \rightarrow 3\pi_D^0$. These columns are normalized so that the events that pass all cuts correspond to the final estimates of background from those sources.

theory calculation implemented in [Gabbiani 97]. Finally, figure 94 shows the final state invariant mass distribution. The remaining background is dominated by $K_L \rightarrow 2\pi_D^0$ background in the low mass sideband, due to energy loss by the electrons.

8.2.6 $K_L \rightarrow \pi^0 e^+ e^- \gamma$ as a background to $K_L \rightarrow \pi^0 e^+ e^-$

Some interest has been raised in the possibility of the rare decay $K_L \rightarrow \pi^0 e^+ e^- \gamma$ being a background to $K_L \rightarrow \pi^0 e^+ e^-$ searches. A monte carlo sample of $K_L \rightarrow \pi^0 e^+ e^- \gamma$




Figure 93. The $e^+e^-\gamma$ invariant mass distribution of the final $K_L \rightarrow \pi^0 e^+ e^- \gamma$ sample is plotted above (crosses) on top of a plot in signal plus data monte carlo (solid histogram). The shaded histogram shows the background only monte carlo. The normalization is absolute.

equivalent to 50 times the KTeV dataset (using the above branching ratio) shows no events faking a $K_L \rightarrow \pi^0 e^+ e^-$ event in the KTeV $K_L \rightarrow \pi^0 e^+ e^-$ analysis. Thus, $K_L \rightarrow \pi^0 e^+ e^- \gamma$ does not present a problem at the current level of the search at $O(10^{-10})$. This is because in order to fake a $K_L \rightarrow \pi^0 e^+ e^-$ event, a photon must be lost from the $K_L \rightarrow \pi^0 e^+ e^- \gamma$ final state. If a photon from the π^0 is lost, then the final state does not reconstruct a π^0 . If the direct photon is lost, then the final state will not satisfy mass or P_t cuts since this photon tends to be energetic.




Figure 94. The $e^+e^-\gamma\gamma$ invariant mass distribution of the final $K_L \rightarrow \pi^0 e^+ e^- \gamma$ sample is plotted above (crosses) on top of a plot in signal plus data monte carlo (solid histogram). The shaded histogram shows the background only monte carlo. The normalization is absolute.

CHAPTER 9

Conclusion

Based on 45 events seen in the the KTeV/E799 data and 3.5 events projected background, we have measured

$$BR(K_L \rightarrow \pi^0 e^+ e^- \gamma) = (2.42 \pm 0.38_{\text{STAT}} \pm 0.11_{\text{SYS}}) \times 10^{-8}. \quad (9.1)$$

This value agrees very well with the calculation given in [Donoghue 97] and [Gabbiani 97] at $O(p^6)$ in Chiral Perturbation Theory. The prediction for the branching ratio is

$$BR(K_L \rightarrow \pi^0 e^+ e^- \gamma)_{O(p^6)} = 2.4 \times 10^{-8} \quad (9.2)$$

at $O(p^6)$ of chiral perturbation theory. At $O(p^4)$, the rate is given by

$$BR(K_L \rightarrow \pi^0 e^+ e^- \gamma)_{O(p^4)} = 1.0 \times 10^{-8} \quad (9.3)$$

The $O(p^6)$ calculation includes vector meson (VMD) exchange effects parameterized by a_v . The value of a_v used in the calculation was -0.96. The measurement is not consistent with the $O(p^4)$ result, given in chapter 1.

Furthermore, the $e^+ e^- \gamma$ invariant mass spectrum agrees well with the monte carlo prediction using the $O(p^6)$ model with VMD effects, and it agrees well with the $\gamma\gamma$ invariant mass spectrum in $K_L \rightarrow \pi^0 \gamma\gamma$ decay, measured by KTeV/E832. A comparison of the two spectra appears in figure 95.

Finally, the rare decay $K_L \rightarrow \pi^0 e^+ e^- \gamma$ does not appear to be a serious background issue in the search for $K_L \rightarrow \pi^0 e^+ e^-$ at current levels. Currently, KTeV/E799 expects to




Figure 95. The shaded histogram shows the $\gamma\gamma$ invariant mass spectrum from the latest KTeV measurement of $K_L \rightarrow \pi^0 \gamma\gamma$ [Alavi 99b]. Overlaid with error bars is the $e^+e^-\gamma$ invariant mass spectrum from the $K_L \rightarrow \pi^0 e^+e^-\gamma$ analysis. The agreement between the two spectra is very good. The cut on (non- π^0) $\gamma\gamma$ mass in the $K_L \rightarrow \pi^0 \gamma\gamma$ analysis is tighter than the corresponding cut on $e^+e^-\gamma$ mass in the $K_L \rightarrow \pi^0 e^+e^-\gamma$ analysis. In the $K_L \rightarrow \pi^0 \gamma\gamma$ analysis, events with (non- π^0) $\gamma\gamma$ mass between 110 MeV/ c^2 and 160 MeV/ c^2 . *K_L → π⁰γγ plot courtesy of S. Taegar.*

set a limit on $\text{BR}(K_L \rightarrow \pi^0 e^+e^-)$ on the order of a few $\times 10^{-10}$. In a $50\times$ KTeV monte carlo sample of $K_L \rightarrow \pi^0 e^+e^-\gamma$ no events were found to fake $K_L \rightarrow \pi^0 e^+e^-$.

APPENDIX A

Note on Error Estimates with Weighted Monte Carlo

A.1 Working with Weighted Monte Carlo

The technique of Monte Carlo integration is often employed when the integral of a function is not easily performed by hand nor easily computed numerically. This often occurs when the integrand is difficult or impossible to integrate in closed form or when the volume of integration is not easily specified [Rubinstein 81],[Kalos 86]. In particle physics, this is often the case with the calculation of the acceptance efficiency for particle decays. In the Monte Carlo method, random points are generated within the integration volume and are then summed to estimate the integral. In general, as with any statistical method of estimation, the errors in the estimation of the integrals are calculable from the distribution of the generated points. It is the purpose of this appendix to derive expressions for these errors when the generated points have weights associated with them.

A.1.1 Errors with Weighted Monte Carlo

We consider a set of elements $x_i \in X$ over which it is desired to perform a Monte Carlo integration. We assume that there exists a weighting function $w_i : X \rightarrow \Re$ and a function to be integrated f . In the following, we will assume that X is a set of generated events and that f is some acceptance function f_A so that $f_A : X \rightarrow \{0,1\}$. For the purposes of this note, the points x_i are equivalent to the ordered pair (w_i, f_{Ai}) . The Monte Carlo integral of the acceptance is the expectation value of f_A over X given the weights, w_i

$$A \equiv \mu = \frac{\sum w_i f_{Ai}}{\sum w_i}. \quad (\text{A.1})$$

The variance of the f_{Ai} is given by the expression

$$\sigma^2 \simeq \frac{\sum w_i (f_{Ai} - A)^2}{\sum w_i}. \quad (\text{A.2})$$

Expression A.2 has not been corrected for the number of degrees of freedom left when the mean A has been estimated from the same set of data X . A way to estimate the number of degrees of freedom for weighted Monte Carlo will be given later. The variance in the mean σ_A^2 is defined as the sum of the products of the individual squared deviates, δ_i^2 , with the squares of the effects each deviation has on the result A , δA_i^2 ,

$$\sigma_A^2 = \sum \delta_i^2 \delta A_i^2, \quad (\text{A.3})$$

where $\delta_i = (f_{Ai} - A)$ and

$$\delta A_i = \frac{w_i f_{Ai}}{\sum w_i} \Big|_{f_{Ai}=0}^{f_{Ai}=1} = \frac{w_i}{\sum w_i}. \quad (\text{A.4})$$

Substituting these expressions into equation A.3, we obtain

$$\sigma_A^2 = \frac{\sum w_i^2 (f_{Ai} - A)^2}{(\sum w_i)^2} = \frac{\sum w_i^2 f_{Ai} - 2A \sum w_i^2 f_{Ai} + A^2 \sum w_i^2}{(\sum w_i)^2}, \quad (\text{A.5})$$

where the fact that f_{Ai} is either 0 or 1 has been used to reduce f_{Ai}^2 to f_{Ai} . Expression A.5 cannot in general be simplified any further; however, a simplified expression results if we consider the case where the mean A is small,

$$A \ll 1. \quad (\text{A.6})$$

In this approximation, the second term in expression A.5 is obviously negligible with respect to the first term. The last term in expression A.5 can be estimated by assuming that

$$\sum w_i^2 \simeq \frac{\sum w_i^2 f_{Ai}}{A}, \quad (\text{A.7})$$

and is therefore approximately the same size as the second term in equation A.5. This assumption may not be valid in cases where there is some devilish correlation between the acceptances, f_{Ai} , and the weights, w_i , so care must be taken when using the approximate form. In particular, it can be shown with a simple "pencil and envelope" test of about 10 points that there exist pathological cases where the assumption in A.7 breaks down. (For a concrete example of this, please refer to section A.4.) In any case, it is worthwhile to explore the conditions under which only the first term survives since it leads directly to the expression for errors in weighted Monte Carlo that has been used before in KTeV and that appears in the literature [Lyons 86]. The principal benefit that accrues when the second two terms of equation A.5 are dropped is that the fractional error in the mean, $\frac{\sigma_A}{A}$, then depends only on the weights of elements, x_i , that satisfy $f_{Ai} = 1$ (ie- the "observed" events,) so that there is no need to keep track of the weights of the "unobserved" events. Keeping only the first term in equation A.5 leads to the approximate expression

$$\sigma_A^2 \simeq \frac{\sum w_i^2 f_{Ai}}{(\sum w_i)^2} \quad (\text{A.8})$$

for the error in the determination of A . The fractional error is given by the expression

$$\frac{\sigma_A^2}{A} \simeq \frac{\sum w_i^2 f_{Ai}}{(\sum w_i f_{Ai})^2}, \quad (\text{A.9})$$

using equation A.1. This is the expression that was used in [Nakaya 94] for the errors in weighted Monte Carlo for the decay $K_L \rightarrow e^+ e^- \gamma \gamma$ and in the book [Lyons 86] for the case of extracting parameters from data and associated weights from detection inefficiencies in a typical nuclear physics experiment.

A.1.2 Estimate of the Number of Degrees of Freedom

Once an expression has been obtained for the fractional error in A , $\frac{\sigma_A}{A}$, one can obtain an estimate of the effective number of degrees of freedom, N_{eff} , with the expression [Lyons 86].

$$\frac{\sigma_A^2}{A^2} \equiv \frac{\sigma_{N_{eff}}^2}{N_{eff}^2} = \frac{1}{N_{eff}}, \quad (\text{A.10})$$

so that using equations A.5 and A.10 and A.1, we obtain

$$N_{eff} = \frac{(\sum w_i f_{Ai})^2}{\sum w_i^2 f_{Ai} - 2A \sum w_i^2 f_{Ai} + a^2 \sum w_i^2} \quad (\text{A.11})$$

Expression A.11 can be used to correct the error obtained for A when there are fewer degrees of freedom available if A and σ_A were computed from the same set of points $X = \{x_i : (w_i, f_{Ai})\}$. For example, the expression for error in the calculation of the mean, σ_A , becomes modified, so that

$$\sigma_A \rightarrow \sigma_A^{N-1} = \sigma_A \times \frac{N_{eff}}{N_{eff} - 1}. \quad (\text{A.12})$$

Note that if we use the approximate expression for the fractional error $\frac{\sigma_A}{A}$ given in equation A.9 above, then

$$N_{eff} = \frac{(\sum w_i f_{Ai})^2}{\sum w_i^2 f_{Ai}}. \quad (\text{A.13})$$

This last approximate expression for N_{eff} is the one given in [Nakaya 94] and [Lyons 86].

A.2 Test of the Error Estimates

To assess the validity of the above formulas in a real setting, a test was set up using both KTeV Monte Carlo for the $K_L \rightarrow e^+e^-\gamma\gamma$ decay mode and for a simple random number Monte Carlo. The former case tests the error formulas in the KTeV experiment for the calculation of acceptance of $K_L \rightarrow e^+e^-\gamma\gamma$, and the latter case explores the validity of the formulas at higher acceptances and with different correlations between the w_i and f_{Ai} . In the following, the points x_i are assumed to be events with event weight given by w_i . If the event is accepted by the analysis, then f_{Ai} is taken to be 1. If the event is not accepted by the analysis, then f_{Ai} is taken to be 0.

A.2.1 Errors with Weighted KTeV Monte Carlo

For this test, a large sample of $K_L \rightarrow e^+e^-\gamma\gamma$ Monte Carlo was generated. This sample had 1.27×10^8 events generated with a generated weight of 4.91×10^5 and a generated sum of squares of weights of 5.26×10^4 . In an analysis of $K_L \rightarrow e^+e^-\gamma\gamma$ by the author, the accepted weight of $K_L \rightarrow e^+e^-\gamma\gamma$ is 2951.0 and the accepted sum of squared weights is 348.3, so that the integrated acceptance of the whole sample is $A_{tot} = (0.601 \pm 0.004)\%$ using the above expressions A.5 and A.8. This large sample was then broken up into 50 smaller subsamples of equal size and the acceptance A_j was recalculated for each subsample. The errors in the acceptance are calculated for each subsample. The general expression A.5 is used to calculate $\sigma_{A_j}^{A.5}$ and then the approximate expression A.8 is used to calculate $\sigma_{A_j}^{A.8}$. Finally, a χ^2 is calculated for each method of calculating error and a comparison is made.

$$\chi_{A.5,A.8}^2 = \sum_{j=1}^{50} \frac{(A_j - A_{tot})^2}{(\sigma_{A_j}^{A.5,A.8})^2}. \quad (\text{A.14})$$

The results are shown in table 26.

In the table, it can be seen that the approximate expression A.8 and the general expression A.5 give very nearly identical results for the errors. The χ^2 statistics also give results of about 1 per degree of freedom indicating that the two formulas are giving error bars of the correct size. Furthermore, one can see that the last two terms in equation A.5

Equation Number	Average Error	χ^2 50 df	Average 1st Term	Average 2nd Term	Average 3rd Term
A.5	2.70×10^{-4}	1.02	7.2×10^{-8}	8.8×10^{-10}	4.1×10^{-10}
A.8	2.69×10^{-4}	1.03	7.2×10^{-8}	N/A	N/A

Table 26. Errors for calculated acceptances using KTeV $K_L \rightarrow e^+e^-\gamma\gamma$ Monte Carlo. 50 trial runs were performed to measure the acceptance. The errors were calculated for each trial run and compared to the observed scatter of individual means A_j about the global mean A_{tot} . Listed in the table are the error expression used to calculate the errors in each trial run, the simple average error calculated in the trial runs, and the χ^2 statistic for the 50 trial runs compared to the global average computed following equation A.14. Also listed at the end of the table are the average sizes of the corresponding terms in expression A.5. In the case of KTeV $K_L \rightarrow e^+e^-\gamma\gamma$ Monte Carlo, the last two terms of expression A.5 are much smaller than the first term.

are indeed negligible with respect to the first term in this case. This is true since the acceptance satisfies approximation A.6. Evidently, the assumption in equation A.7 holds in this case as well.

A.2.2 Results with Simple Random Number Monte Carlo

In this test, a set of events $x_i = (w_i, f_{Ai})$ are generated using a simple random number generator. The w_i are generated uniformly over the interval $[0, 1]$ and the f_{Ai} are generated for three different cases. Recall that the f_{Ai} are either 0 or 1 depending on whether or not the event is accepted. Let the fraction of elements x_i at a given w_i with $f_{Ai} = 1$ be denoted by $F^1(w_i)$. In the first case, the acceptance is chosen to be uncorrelated with the w_i ; the f_{Ai} are thrown such that $F^1(w_i) = C_1$ where C_1 is some constant between 0 and 1. In the second case, the acceptance is chosen to be positively correlated with the w_i ; the f_{Ai} are thrown such that $F^1(w_i) = C_2 w_i$. Finally, in the third case the acceptance is chosen to be anti-correlated with the w_i ; the f_{Ai} are thrown such that $F^1(w_i) = C_3(1 - w_i)$. One can see that in the first case, the chance of an event being accepted is unrelated to its weight. In the second case, the events with large weights tend to have more acceptance. Finally, in the third case, the events with small weights have larger acceptance. In each

A	Uncorrelated		Correlated		Anti-correlated	
	$\chi^2_{A.5}$ 100 df	$\chi^2_{A.8}$ 100 df	$\chi^2_{A.5}$ 100 df	$\chi^2_{A.8}$ 100 df	$\chi^2_{A.5}$ 100 df	$\chi^2_{A.8}$ 100 df
0.01	1.39	1.38	1.56	1.54	1.21	1.20
0.05	0.91	0.86	1.08	1.03	1.01	0.98
0.1	1.19	1.07	1.28	1.15	1.05	0.98
0.3	1.19	0.83	1.11	0.74	1.15	0.93
0.5	1.17	0.59	1.47	0.66	1.03	0.68
0.75	1.34	0.34	1.30	0.30	0.83	0.28

Table 27. Errors for calculated acceptances using simple randomly generated weights and acceptances. Each number in the table represents the χ^2 per degree of freedom fit of the averages of a set of 100 trials of 5000 events each to the global average. A low χ^2 indicates that the errors are overestimated in that set of trials. The χ^2 was calculated using both equation A.5 and equation A.8 for three different cases : acceptance is uncorrelated with weight, acceptance is correlated with weight, and acceptance is anti-correlated with weight. The normalization of the acceptance is then varied so that the average acceptance varies from 1% to 75%.

case, the value of the constant was varied to maintain the same value of the overall mean acceptance A.

In each case and for each chosen value of the constant C, 5000 events x_i were generated in 100 trial runs for a total of 500000 events. From the values A_j , the mean in each run, and the errors σ_{A_j} , and the mean over all 500000 A_{tot} , χ^2 statistics were generated following equation A.14. The χ^2 is then studied as a function of mean A_{tot} , the method used to calculate the error (either equation A.5 or equation A.8,) and the different assumptions about the correlation between the f_{Ai} and the w_i . These results are shown in table 27.

From table 27 one can see that the approximate expression for the error A.8 breaks down giving errors that are too large when the acceptance is high. The effect seems to be largely independent of the particular assumptions made here about the correlations with the weights. In figure 96, the ratio of the χ^2 statistics in the general case to those in the approximate case is shown for each assumption about the correlation as a function of




Figure 96. The ratio of χ^2 calculated using the approximate formula for the errors A.8 to that calculated using the general formula A.5. The ratios for the uncorrelated assumption are the circles, the ratios for the correlated assumption are the squares, and the ratios for the anti-correlated case are the triangles. Since the general formula A.5 is seen to give χ^2 results of about 1 per degree of freedom, a ratio of less than one indicates that the approximate formula is producing abnormally large errors.

acceptance. One can get an approximate rule of thumb that the approximated errors are about a factor of $(1 + A)$ larger than the errors calculated using the more general form.

A.3 Conclusion

In this appendix, expressions for errors in the Monte Carlo estimation of integrals were given for the case of weighted Monte Carlo. It was found that the general expression for the error in the Monte Carlo estimation of an acceptance function f_{Ai} over a set of events x_i with weights w_i is given by the formula

$$\frac{\sigma_A^2}{A^2} = \frac{\sum w_i^2 (f_{Ai} - A)^2}{(\sum w_i f_{Ai})^2} = \frac{\sum w_i^2 f_{Ai}^2 - 2A \sum w_i^2 f_{Ai} + A^2 \sum w_i^2}{(\sum w_i f_{Ai})^2}.$$

In the case where the acceptance is small, and we make an additional assumption about the weakness of correlation between the function f_A and the weights w_i , this expression simplifies to

$$\frac{\sigma_A^2}{A^2} \simeq \frac{\sum w_i^2 f_{Ai}^2}{(\sum w_i f_{Ai})^2},$$

i	w_i	f_{Ai}	i	w_i	f_{Ai}
1	0.01	1	6	0.5	0
2	0.01	1	7	0.5	0
3	0.01	1	8	0.5	0
4	0.01	1	9	0.5	0
5	0.01	1	10	0.5	0

Table 28. A small "pencil and envelope" sample of events. The events are characterized solely by their weights and their acceptance bits.

which is what was used in KTeV, and only depends on the "observed" weights. These expressions were tested using weighted KTeV Monte Carlo for the decay $K_L \rightarrow e^+e^-\gamma\gamma$ and it was found that both the exact expression and the approximate expression give correct estimations for the errors. At higher acceptances, as tested by a simple random number weighted Monte Carlo, the approximate expression gives errors that are too large as gauged by the χ^2 statistic. The errors are too large typically by a factor of $(1 + A)$ for the cases explored in this note.

A.4 A Small Example

The general expression for error when using weighted monte carlo, given by equation A.5, can be simplified to give the expression A.8 when the conditions A.6 and A.7 are satisfied. The case where condition A.6 fails was explored in the note, and the case where condition A.6 holds but condition A.7 fails will be explored here. Consider the set of events presented in table 28.

In this case, the mean A works out to 0.0196, while the error calculated using general expression A.5 is ± 0.012 and the error calculated using the approximate expression A.8 is ± 0.008 . In this case, the approximate formula gives errors that are different from the exact case.

APPENDIX B

Confidence Intervals With Background

Error :

An Extension to the Method of Feldman
and Cousins

The Feldman and Cousins [Feldman 98] method of constructing confidence intervals uses the construction of Neyman [Neyman 37] to build confidence belts while proposing a solution to ambiguities inherent in the construction when applying it to this problem. However, the method of Feldman and Cousins as published does not contain within it a prescription for propagating errors in the calculation of background. This appendix proposes a method for incorporating such errors.

B.1 The Method of Feldman and Cousins in Brief

The method of Feldman and Cousins to find confidence intervals in the analysis of small signals in the presence of background rests upon the construction of confidence belts. The discussion in this section relies heavily on the excellent discussion found in [Feldman 98]. Consider an experiment measuring some Poisson-like parameter μ in the presence of exact background b . A measurement of μ will yield an observation n , where n is an integer. Also desired is to quote a confidence interval with confidence α . In order to calculate the intervals, one can construct a confidence belt. To do this, for each value of μ , one finds a set S of n such that

$$\sum_{n \in S} P(n, \mu; b) \geq \alpha, \quad (\text{B.1})$$

where

$$P(n, \mu; b) = \frac{(\mu + b)^n e^{-\mu-b}}{n!} \quad (\text{B.2})$$

is the Poisson probability function. It is understood that we want the sum to be as close to α as possible; but because of the discrete nature of the sum, we will overshoot from time to time. This property is known as *conservatism* and is generally undesirable. (The sum should represent the stated confidence level !) Usually, we take the set S as containing n that define a simply connected interval, called by Feldman and Cousins the *acceptance region*. This procedure is repeated for all μ . When plotted in the $\mu - n$ plane, the acceptance regions form a jagged structure known as the confidence belt. The confidence interval at the desired level α for a given observation n is then simply given

99/01/10 17.13




Figure 97. A confidence belt constructed using the method of Feldman and Cousins in the $\mu - n$ plane. The belt constructed above is for the case of an exact background of 3.0 at a confidence level of 90%. The horizontal intervals bounded by the jagged curves are known as acceptance intervals; the n included in each acceptance interval satisfy relation B.1. A confidence interval at a specific n is the smallest interval in μ that contains all of the acceptance regions intersected by a vertical line drawn at n .

by the extrema in μ of the set of acceptance regions that intersect a vertical line drawn through the confidence belt at n . See figure 97 for a pictorial depiction of this method.

There is an ambiguity, however, in the construction as I have stated it. I have not specified *which* n to include in each set S of acceptance regions. One popular choice is to include n on the low side which leads directly to *upper confidence limits*. Another

is to require summed probability outside of each acceptance region on either side of the acceptance region to balance, leading to *central confidence intervals*. However, difficulties arise with either of these choices that are not *a priori* obvious. First, if the experimenter lets the data suggest which prescription to use, then the resulting intervals may actually *undercover* the stated confidence level. This behavior, observed among experimenters, has been dubbed *flip-flopping* by Feldman and Cousins. Secondly, with either of the above constructions, if the experimenter observe $n = 0$ events, the resulting confidence interval is the empty set in the Poisson case.

The method suggested by Feldman and Cousins introduces a new prescription for including n in the set S . Namely, a best value μ_b is determined for each n as

$$\mu_b = \max(0, n - b). \quad (\text{B.3})$$

Then the n are ranked in decreasing order of the ratio

$$R = \frac{P(n, \mu; b)}{P(n, \mu_b; b)}. \quad (\text{B.4})$$

n are simply included in the set S in order of descending R until the desired confidence is reached (or exceeded). The wonderful thing about this prescription is that it gives a non-empty confidence interval for the case $n = 0$ and it eliminates the need for flip-flopping. It gives an upper confidence limit in the regions of the confidence belt where appropriate and gives a two sided interval otherwise. However, the method as presented in [Feldman 98] only deals with exact backgrounds.

B.2 An Extension to the Method of Feldman and Cousins

My extension to the method of Feldman and Cousins involves error in the determination of b . In order to make the discussion more concrete, imagine that the background is distributed according to some probability density function of its own, $P_b(b')$. In order to take full advantage of the Neyman construction, this new probability density function is inserted into the calculation at an early stage during the construction of acceptance regions. Instead of using equation B.1 to define an acceptance region, I use the following

requirement on the set S' , the new acceptance region :

$$\sum_{n \in S'} \int_0^\infty db' P_b(b') P(n, \mu; b') \geq \alpha. \quad (\text{B.5})$$

When the integral in B.5 is intractable, (which is often the case,) one can replace the integral with a discrete sum over some reasonable sampling region in b' .

$$\sum_{n \in S'} \sum_{b' \in \mathcal{B}} P_b(b') P(n, \mu; b') \geq \alpha. \quad (\text{B.6})$$

where \mathcal{B} is some set of background points considered as representative of the interval. Using this expression for the numerator in the ratio B.4, the n are re-ordered for inclusion into the set S' to define a new acceptance region.

$$R^{new} = \frac{\sum_{b' \in \mathcal{B}} P_b(b') P(n, \mu; b')}{P(n, \mu_b; b)}. \quad (\text{B.7})$$

where sadists can replace the sum by an integral if desired, and the plain b now refers to the mean b . The construction of the full confidence belt and calculation of the confidence intervals then proceeds as before. The new confidence belt is smeared primarily in regions where the error in b is of order n and μ , leading to wider confidence intervals there.

In the calculation of confidence intervals in the case of exact background, it should be noted that Feldman and Cousins find several minor anomalies. For example, it sometimes happens that the set of acceptance intervals intersected by a vertical line at n in the $\mu - n$ plane is not simply connected. Following Feldman and Cousins, the confidence interval is taken to be the minimum interval in μ that contains all of the intersecting acceptance regions. This anomaly arises because of the discrete nature of the Poisson probability density function. The fix for this anomaly, alas, adds to the conservatism of the interval. Other anomalies are also addressed in [Feldman 98]. Similarly, a new anomaly is introduced in the calculation of confidence intervals with error. This new anomaly consists of the observation that some acceptance regions actually *shorten* in a small number of cases. The anomaly is the result of a correction to the natural conservatism of some of the acceptance regions when calculated in the exact background case. When this happens, the anomaly can be addressed by taking the acceptance interval to be the union of the points in the exact case acceptance region, S , with the error background case acceptance region,

S' . This last fix also adds a bit to the overall conservatism of the final confidence interval; however, it also maintains a consistency of conservatism with the exact background case, considered the default.

B.3 Short Table of Confidence Intervals With Errors in Background

In this section, a short table is given of new confidence intervals calculated using the method of Feldman and Cousins with the extension to handle errors in the background. The background sampling distribution in each case is Gaussian. The confidence level is 90%.

N	0.0	0.5	1.0	1.5	2.0	2.5	3.0	Error
0	0.00,2.44	0.00,1.94	0.00,1.61	0.00,1.33	0.00,1.26	0.00,1.18	0.00,1.08	exact
		0.00,1.94	0.00,1.61	0.00,1.33	0.00,1.28	0.00,1.18	0.00,1.09	10%
		0.00,1.95	0.00,1.61	0.00,1.33	0.00,1.37	0.00,1.18	0.00,1.15	20%
1	0.11,4.36	0.00,3.86	0.00,3.36	0.00,2.91	0.00,2.53	0.00,2.19	0.00,1.88	exact
		0.00,3.86	0.00,3.36	0.00,2.91	0.00,2.53	0.00,2.19	0.00,1.88	10%
		0.00,3.86	0.00,3.36	0.00,2.91	0.00,2.53	0.00,2.19	0.00,1.88	20%
2	0.53,5.91	0.03,5.41	0.00,4.91	0.00,4.41	0.00,3.91	0.00,3.45	0.00,3.04	exact
		0.03,5.41	0.00,4.91	0.00,4.41	0.00,3.91	0.00,3.45	0.00,3.04	10%
		0.03,5.41	0.00,4.91	0.00,4.41	0.00,3.92	0.00,3.45	0.00,3.04	20%
3	1.10,7.42	0.60,6.92	0.10,6.42	0.00,5.92	0.00,5.42	0.00,4.92	0.00,4.42	exact
		0.60,6.92	0.10,6.42	0.00,5.92	0.00,5.42	0.00,4.92	0.00,4.43	10%
		0.60,6.92	0.00,6.42	0.00,5.92	0.00,5.42	0.00,4.92	0.00,4.43	20%
4	1.47,8.60	1.17,8.10	0.74,7.60	0.25,7.10	0.00,6.60	0.00,6.10	0.00,5.60	exact
		1.17,8.10	0.74,7.60	0.24,7.12	0.00,6.64	0.00,6.16	0.00,5.70	10%
		1.17,8.11	0.73,7.64	0.21,7.22	0.00,6.86	0.00,6.42	0.00,5.92	20%
5	1.83,9.99	1.53,9.49	1.25,8.99	0.93,8.49	0.43,7.99	0.00,7.49	0.00,6.99	exact
		1.53,9.49	1.25,8.99	0.93,8.49	0.43,7.99	0.00,7.49	0.00,6.99	10%
		1.53,9.49	1.24,8.99	0.90,8.49	0.38,7.99	0.00,7.49	0.00,6.99	20%

Table 29. 90% confidence Intervals calculated using the method of Feldman and Cousins with exact errors and with 10% and 20% Gaussian errors in the background. The cases with Gaussian error were each calculated on a grid of 11 points $\pm 5\sigma$ about the stated mean background. Most limits do not change, but many change by up to a few percent.

APPENDIX C

Multi Particle Kinematics in Two E-Z Pieces

Analysis of purely kinematical variables is often helpful to separate signal from background. However, the particular variables used tend to vary from analysis to analysis. In this appendix, I present a method which may be of general use when the background under consideration is a fragment of some real decay or mixed up with accidental activity.

Suppose we have an unstable parent particle of mass M that decays into $n > 2$ daughter particles. By grouping the decay products into two groups, we can perform the following purely kinematical analysis. Let m_1 and m_2 be the invariant masses of the two daughter systems respectively. Then in the parent Center of Mass (C.M.) frame, the energy and momentum of the first daughter system is given by

$$E_1^* = \frac{1}{2M}(M^2 - m_2^2 + m_1^2) \quad (\text{C.1})$$

$$P_1^* = \frac{1}{2M}\lambda^{\frac{1}{2}}(M^2, m_1^2, m_2^2) \quad (\text{C.2})$$

where $\lambda(a, b, c) = a^2 + b^2 + c^2 - 2(ab + bc + ac)$.

If it is possible to appeal to measurement to obtain the transverse component of momentum of the first daughter system with respect to the parent particle, P_{T1} ,¹ boost direction in the lab frame, then we can solve for the square of the component of momentum parallel to the boost direction in the parent C.M. frame.

$$P_{1,\parallel}^{*2} = \frac{1}{4}\lambda(M^2, m_1^2, m_2^2) - P_{T1}^2 \quad (\text{C.3})$$

At this point, it is worth noting that this quantity corresponds closely to a variable in the literature called *pp0kin*. Here, I will use \mathcal{P}_1 to refer to our analogue of pp0kin.

$$\mathcal{P}_1 = \frac{P_{1,\parallel}^{*2}}{P_{T1}^2 + m_1^2} \quad (\text{C.4})$$

For a well measured decay process, \mathcal{P}_1 must be non-negative. For example, in the decay process $K \rightarrow \pi^\pm e^\mp \nu$, we can let the first grouping represent the charged decay products and the second grouping represent the neutrino. $K \rightarrow \pi^\pm e^\mp \nu$ events are selected over the also copious $K_L \rightarrow \pi^+ \pi^- \pi^0$ decays for our CsI and TRD calibration samples.

¹Note : P_{Ti} here is distinct from the usual P_T defined in many kaon analyses. The usual P_T is the transverse momentum of the entire event with respect to the initial kaon direction and thus provides a measure of event quality.

In decays where we observe all of the mass, then there exists a definite relationship between \mathcal{P}_1 and m_2 implied by equations C.4 and C.3. That is, if we plot events from whole decays in the $\mathcal{P}_1 - m_2$ plane, then the events will fall along a kinematic boundary. For example, in $K_L \rightarrow 2\pi_D^0$ decay, let the first grouping be the e^+e^- pair and the second grouping be the $\gamma\gamma\gamma$ system. Figure 98 shows the population of simulated $K_L \rightarrow 2\pi_D^0$ events on the $\mathcal{P}_1 - m_2$ plane. On the other hand, consider fragments from the decay $K_L \rightarrow 3\pi_D^0$ in which two photons are lost. The kinds of particles in the final state are thus the same as those in the complete $K_L \rightarrow 2\pi_D^0$ final state. However, because the relations in equations C.1 and C.2 are no longer satisfied (under the assumption that M is the same for both the complete $K_L \rightarrow 2\pi_D^0$ event and the fragmentary $K_L \rightarrow 3\pi_D^0$ event), then the equalities in C.3 and C.4 no longer hold. In short, the fragmentary events will populate a broad region in the $\mathcal{P}_1 - m_2$ plane bordered by the exact relation implied by the same equations. Figure 98 shows the population of simulated $K_L \rightarrow 3\pi_D^0$ events on the $\mathcal{P}_1 - m_2$ plane.

Yet all is not perfect for it turns out that fragmentary events for which equations C.1 and C.2 hold in the approximate, the events will populate a region close to the border. In other words, for fragmentary events in which the energy lost is small compared to the energy in the whole event, then these events tend to overlap with complete events in the $\mathcal{P}_1 - m_2$ plane. Unfortunately, these are the events which also pass mass cuts and event quality cuts. Yet the $\mathcal{P}_1 - m_2$ cut remains a very effective tool for rooting out backgrounds in mass sidebands, backgrounds with accidental energy, and backgrounds in cases when it is undesirable to use the more standard mass and event quality cuts. Table X shows the effectiveness of the $\mathcal{P}_1 - m_2$ cut against $K_L \rightarrow 3\pi_D^0$ in a $K_L \rightarrow 2\pi_D^0$ analysis in various mass sidebands just below the kaon mass.

Finally, in order to more easily optimize a cut in the $\mathcal{P}_1 - m_2$ plane, it is suggested to induce a curved cut in this plane that approximates the shape of the kinematical boundary. This is easily done by cutting on some linear combination of $\log \mathcal{P}$ and m_2 that transforms the kinematical boundary region into a region that approximates a straight line. One simply cuts on the projection of this transformed variable to induce a curved cut in the original plane. Such a variable, V_c , in the case of $K_L \rightarrow 2\pi_D^0$ and $K_L \rightarrow 3\pi_D^0$

99/01/10 17.04




Figure 98. The distributions of \mathcal{P}_1 versus m_2 for $K_L \rightarrow 3\pi_D^0$ and $K_L \rightarrow 2\pi_D^0$ Monte Carlo after requiring two charged tracks and 5 electro-magnetic clusters in the CsI calorimeter. The $K_L \rightarrow 3\pi_D^0$ events (top) are fragmentary and thus occupy an extended region in the $\mathcal{P}_1 - m_2$ plane. The $K_L \rightarrow 2\pi_D^0$ events are complete and fall along a kinematic boundary in the $\mathcal{P}_1 - m_2$ plane. A small fraction of $K_L \rightarrow 2\pi_D^0$ (bottom) events do in fact populate nether regions of the plot, but these are due to misreconstructed events or events with accidental energy. These poor events would fail to satisfy more stringent analysis cuts anyway.




Figure 99. The top plot shows the $K_L \rightarrow 3\pi_D^0$ and the $K_L \rightarrow 2\pi_D^0$ distributions in the $\log \mathcal{P}_1 - m_2$ plane. A line is drawn parallel to the cut in this plane that gives the best separation between the two distributions. Requiring in an analysis that the events be located above or below this line is equivalent to cutting on the variable $V_c \equiv \log \mathcal{P}_1 - 15m_2$ where the line in the plot corresponds to $V_c = 0$. One can plot the distribution of events transverse to this defining line. Such a plot is shown in the bottom plot where the same Monte Carlo distributions of $K_L \rightarrow 3\pi_D^0$ and $K_L \rightarrow 2\pi_D^0$ are overlaid.

discrimination is

$$V_c = \log \mathcal{P}_1 - 15.4m_2 + 4.6 \quad (\text{C.5})$$

For an illustration of this cut, see figure 99. Also, the distribution of events lying above or below V_c is shown in figure 81 (bottom).

REFERENCES

- [Aherton 80] H.W. Aherton et al., CERN Preprint 80-07, CERN (1980)
- [Alavi 99a] A. Alavi-Harati, Ph.D. Thesis, University of Wisconsin, Madison, (1999)
- [Alavi 99b] A. Alavi-Harati, et al. (The KTeV Collaboration) hepex-9902029 (to appear in Phys. Rev. Lett.)
- [Angelopoulos 98] A. Angelopoulos, et al. (CPLEAR collaboration) Phys. Lett. B444, 43 (1998)
- [Artru 75] X. Artru, et al. Phys. Rev. D12,1289 (1975)
- [Barker 90] A. Barker, et al. Phys. Rev. D41, 3546 (1990)
- [Barr 2] G.D. Barr, et al. Phys. Lett. B284, 440 (1992)
- [Behrend 91] H. Behrend et al. Z. Phys. C49, 401 (1991)
- [Belusevic 96] R. Belusevic., Neutral Kaons, KEK preprint 96-182, (1996)
- [Bielajew 91] A.F. Bielajew. Preprint PIRS-0287(1991)
- [Bown 96] C. Bown, et al. Nucl. Inst. Meth. A369, 248 (1996)
- [Briere 95] R. Briere, Ph.D. Thesis, University of Chicago, (1995)
- [Carosi 90] R. Carosi, et al., Phys. Lett. B 237, 303 (1990)
- [Carroll 80] A.S. Carroll, et al. Phys. Rev. Lett. 44, 525 (1980)

- [Cherry 74] M. Cherry, et al. Phys. Rev. D10, 3594 (1974)
- [Christenson 64] J.H. Christenson, et al. Phys. Rev. Lett. 13, 138 (1964)
- [Colangelo 98] G. Colangelo and G. Isidori, hep-ph-9808487, (1998)
- [D'Ambrosio 96] G. D'Ambrosio and G. Isidori. Int. J. Mod. Phys. A13, 1 (1998) (hep-ph-9611284)
- [D'Ambrosio 97] G. D'Ambrosio and J. Potoles, Nucl. Phys. B492,417 (also hep-ph-9610244,) (1997)
- [D'Ambrosio 98] G. D'Ambrosio, et al., hep-ph-9808289, (1998)
- [Donoghue 87] J. Donoghue, B. Holstein, and G. Valencia, Phys. Rev. D 35, 2769 (1987)
- [Donoghue 95] J. Donoghue and F. Gabbiani., Phys. Rev. D 51, 2187 (1995)
- [Donoghue 97] J. Donoghue and F. Gabbiani., Phys. Rev. D56, 1605 (1997) (also updated in hep-ph-9702278v3)
- [Durand 75] L. Durand. Phys. Rev. D11, 89 (1975)
- [Ecker 87] G. Ecker, A.Pich, and E. DeRafael. Nucl. Phys. B 291, 692 (1987)
- [Ecker 88] G. Ecker, A.Pich, and E. DeRafael. Nucl. Phys. B 303, 665 (1988)
- [Ecker 90] G. Ecker, A.Pich, and E. DeRafael. Phys. Lett. B237, 481 (1990)
- [Ecker 93] G. Ecker, A.Pich, and E. DeRafael. Phys. Lett. B304, 347 (1993)
- [Feldman 98] G. Feldman and R. Cousins. Phys. Rev. D 57, 3873 (1998)
- [Gabbiani 97] F. Gabbiani. ADS.f : A Fortran Program to Calculate $\text{BR}(K_L \rightarrow \pi^0 e^+ e^- \gamma)$, Private Communication
- [Gabbiani 97] F. Gabbiani., Ph.D. Thesis, The University of Massachusetts at Amherst, (1997) (also hep-ph-9901262v3)

- [Gabbiani 99] F. Gabbiani. Private Communication.
- [Gell-Mann 55] Gell-Mann, M. and Pais, A. Phys. Rev. 97, 1387 (1955)
- [Gibbons 93a] L. Gibbons, Ph.D. Thesis, University of Chicago, (1993)
- [Gibbons 93b] L.K. Gibbons et al., Phys. Rev. Lett. 61, 2661 (1988)
- [Gibbons 93b] L.K. Gibbons et al., Phys. Rev. Lett. 70, 1199 (1993)
- [Hanagaki 98] K. Hanagaki, Ph.D. Thesis, Osaka University, (1998)
- [Haney 92] Haney, M.J., et al. *Conference Record of the 1992 Nuclear Science Symposium*, p. 338 (1992)
- [Harris 94] D. Harris, Ph.D. Thesis, University of Chicago, (1994)
- [Harris 93] D. Harris, et al. Phys. Rev. Lett., (1993)
- [Heiliger 93] P. Heiliger and L.M. Seghal, Phys. Rev. D 47, 4920 (1993)
- [Isidori 99] G. Isidori. hep-ph-9902235, also, Talk presented at the “International Workshop on CP Violation in K”, 18-19 December 1998, KEK-Tanashi, Tokyo (1999)
- [Kalos 86] M.H. Kalos and P.A. Whitlock, Monte Carlo Methods, Volume 1, John Wiley and Sons, New York (1986)
- [Kambor 91] Kambor, J. et al. Phys. Lett. B261, 496 (1991)
- [Kroll 55] N. Kroll and W. Wada, Phys. Rev. 98, 1355 (1955)
- [KTeVana] KTeV Analysis Code, Version 4.11. Unpublished. (1998)
- [KTeVmc] KTeV Monte Carlo Code, Version 4.11. Unpublished. (1998)
- [KTeV Memo] A KTeV Note, to be found later.
- [KTeV 97] Kobilacik, T., et al. *KTeV Internal Memo 97*

- [KTeV 102] Barker, T. and Nauenberg, U. *KTeV Internal Memo 102*
- [KTeV 113] Schnetzer, S., et al. *KTeV Internal Memo 113*
- [KTeV 124] Jensen, D.A. and Hsiung, Y.B. *KTeV Internal Memo 124*
- [KTeV 132] Barker, T. and Nauenberg, U. *KTeV Internal Memo 132*
- [KTeV 141] Coleman, R. *KTeV Internal Memo 141*
- [KTeV 146] G.E. Graham, et al. *KTeV Internal Memo 146*
- [KTeV 154] M. Daum, et al. *KTeV Internal Memo 154*
- [KTeV 158] Graham, G.E. *KTeV Internal Memo 158*
- [KTeV 177] Hanagaki, K. and Yamanaka, T. *KTeV Internal Memo 177*
- [KTeV 184] O'Dell, V. *KTeV Internal Memo 184*
- [KTeV 185] Y.B. Hsiung, et al. *KTeV Internal Memo 185*
- [KTeV 187] Jennings, J., et al. *KTeV Internal Memo 187*
- [KTeV 191] Schnetzer, S., et al. *KTeV Internal Memo 191*
- [KTeV 196] Kassabian, G., et al. *KTeV Internal Memo 196*
- [KTeV 198] Barker, T. *KTeV Internal Memo 198*
- [KTeV 199] Barker, T. *KTeV Internal Memo 199*
- [KTeV 201] Mikelsons, P. *KTeV Internal Memo 201*
- [KTeV 203] Yarema, R.J., et al. *KTeV Internal Memo 203*
- [KTeV 207] Alexopoulos, T., et al. *KTeV Internal Memo 207*
- [KTeV 213] Coleman, R., et al. *KTeV Internal Memo 213*
- [KTeV 213] Shawhan, P. *KTeV Internal Memo 257*

- [KTeV 213] Arenton, A. *KTeV Internal Memo 265*
- [KTeV 298] Bown, C., et al. *KTeV Internal Memo 298*
- [KTeV 317] Roodman, A. *KTeV Internal Memo 317*
- [KTeV 328] P. Shawhan *KTeV Internal Memo 328*
- [KTeV 361] Kessler, R. *KTeV Internal Memo 361*
- [KTeV 369] Bright, S. *KTeV Internal Memo 369*
- [KTeV 389] O'Dell, V. *KTeV Internal Memo 389*
- [KTeV 403] Roodman, A. *KTeV Internal Memo 403*
- [KTeV 577] Roodman, A. *KTeV Internal Memo 577*
- [KTeV 608] Bellantoni, L. et al. *KTeV Internal Memo 608*
- [Lande 56] E. Lande, et al. Phys. Rev. 103, 1901 (1956)
- [Leo 87] W.R. Leo. Techniques for Nuclear and Particle Physics Experiments, Springer-Verlag, New York (1987)
- [Lyons 86] L. Lyons, Statistics for Nuclear and Particle Physicists, Cambridge University Press, Cambridge, U.K. (1986)
- [Malensek 81] A.J. Malensek, FNAL Preprint FN-341 and errata, FNAL (1981)
- [Marshak 93] R. Marshak. The Conceptual Foundations of Modern Particle Physics, World Scientific, New York, (1993)
- [Nactmann 89] O. Nachtmann., Elementary Particle Physics Concepts and Phenomena, Springer-Verlag, Berlin (1989)
- [Nakaya 94] T. Nakaya, Ph.D. Thesis, Osaka University, Osaka, Japan (1994)
- [Neyman 37] Philos. Trans. R. Soc. London A236, 333 (1937)

- [Papadimitriou 93] V. Papadimitriou, Ph.D. Thesis, University of Chicago, (1993)
- [Papadimitriou 91] V. Papadimitriou, et al. Phys. Rev. D44, 573 (1991)
- [Patterson 92] J.R. Patterson, Ph.D. Thesis, University of Chicago, (1992)
- [PDG 98] Particle Data Group, The Euro. Phys. J C3, Review of Particle Physics, (1998)
- [Ping 94] G. Ping, et al. Phys. Rev. Lett. 72, 3000 (1994)
- [Quinn 99] G.B. Quinn, Ph.D. Thesis, University of Chicago, (1999) (To appear soon.)
- [Rubinstein 81] R. Rubinstein, Simulation and the Monte Carlo Method, John Wiley and Sons, New York (1981)
- [Sachs 87] Sachs, R.G. The Physics of Time Reversal, University of Chicago Press, Chicago (1987)
- [Sauli 77] F. Sauli CERN 77-09, (1977)
- [Schwingenheuer 95] B. Schwingenheuer et al., Phys. Rev. Lett. 74, 4376 (1995)
- [Scott 63] W.T. Scott. Rev. Mod. Phys. 35, 231 (1963)
- [Shawhan 99] P. Shawhan, Ph.D. Thesis, University of Chicago, (1999) (To appear soon.)
- [Storm 70] E. Storm and H. Isreal. Nucl. Data Tables A7, 565-681 (1970)
- [Ter-Mikaelian 72] M.C. Ter-Mikaelian. High Energy Electromagnetic Processes in Condensed Media. Wiley-Interscience, New York (1972)
- [Prasad 98] Valmiki Prasad, private communication.
- [Zimmerman 99] E.D. Zimmerman, Ph.D. Thesis, University of Chicago, (1999)